



Applications

Note

Before using this information, be sure to read the general information under “Notices” on page 1433.

Compilation date: April 14, 2004

© Copyright International Business Machines Corporation 2004. All rights reserved.

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

How to send your comments	xxv
Chapter 1. Welcome to applications	1
Chapter 2. Using Web applications	11
Web applications	11
web.xml file	11
Migrating Web application components	13
Default Application	15
Servlets	16
Developing servlets with WebSphere Application Server extensions	16
Application lifecycle listeners and events	17
Listener classes for servlet context and session changes	17
Example: com.ibm.websphere.DBConnectionListener.java	18
Servlet filtering	18
Filter, FilterChain, FilterConfig classes for servlet filtering	19
Example: com.ibm.websphere.LoggingFilter.java	19
Configuring page list servlet client configurations	20
autoRequestEncoding and autoResponseEncoding	23
Examples: autoRequestEncoding and autoResponseEncoding encoding examples	23
JavaServer Pages files	24
Developing JavaServer Pages files with WebSphere extensions	24
Tag libraries	25
tsx:dbconnect tag JavaServer Pages syntax	25
dbquery tag JavaServer Pages syntax	26
dbmodify tag JavaServer Pages syntax	27
tsx:getProperty tag JavaServer Pages syntax and examples	27
tsx:userid and tsx:passwd tag JavaServer Pages syntax	28
tsx:repeat tag JavaServer Pages syntax	29
Example: Combining tsx:repeat and tsx:getProperty JavaServer Pages tags	29
Example: tsx:dbmodify tag syntax	29
Example: Using tsx:repeat JavaServer Pages tag to iterate over a results set	30
JspBatchCompiler tool	32
Bean Scripting Framework	33
Developing Web applications	33
Disabling JavaServer Pages run-time compilation	34
Example: Converting JavaScript source to the Bean Scripting Framework	35
Scenario: Creating a Bean Scripting Framework application	36
Example: Bean Scripting Framework code example	41
Web modules	44
Assembling Web applications	45
Context parameters	45
Security constraints	46
Servlet mappings	46
Invoker attributes	46
Error pages	46
File serving	46
Initialization parameters	46
Servlet caching	47
Web components	47
Web property extensions	47
Web resource collections	47
Welcome files	47

Troubleshooting tips for Web application deployment	47
Modifying the default Web container configuration	48
Web container	49
Web container settings	49
Web Module Deployment settings	49
Web container custom properties.	50
Web applications: Resources for learning.	51
Chapter 3. Managing HTTP sessions	53
Sessions	53
Migrating HTTP sessions	53
Developing session management in servlets	54
Example: SessionSample.java	55
Assembling so that session data can be shared	56
Session security support	56
Session management support	57
Configuring session management by level	58
Session tracking options	59
Session tracking with cookies	59
Session tracking with URL rewriting	59
Session tracking with SSL information	60
Configuring session tracking	61
Serializing access to session data	61
Session Management settings	61
Cookie settings	64
Configuring session tracking for Wireless Application Protocol (WAP) devices	65
Distributed sessions	65
Session recovery support	65
Distributed Environment settings	66
Configuring for database session persistence	66
Switching to a multirow schema	66
Configuring tablespace and page sizes for DB2 session databases	67
Database settings	67
Multirow schema considerations	68
Memory-to-memory replication.	69
Memory-to-memory topology: Peer-to-peer function with a local replicator	70
Memory-to-memory topology: Peer-to-peer function with remote/isolated replicators	71
Memory-to-memory topology: Client/server function with remote replicators	72
Memory-to-memory topology: Client/server function with isolated replicators	73
Memory-to-memory session partitions	74
Clustered session support	74
Configuring memory-to-memory replication for the peer-to-peer function with a local replicator (default memory-to-memory replication)	75
Memory-to-memory sessions settings	75
Configuring memory-to-memory replication for peer-to-peer functions with remote/isolated replicators	76
Configuring memory-to-memory replication for the client/server function using isolated replicators	77
Configuring memory-to-memory replication for the client/server function using remote replicators	79
Tuning session management	80
Configuring scheduled invalidation	80
Configuring write contents	81
Configuring write frequency	82
Base in-memory session pool size	82
Controlling write operations	83
Tuning parameter settings	83
Tuning parameter custom settings	84
Best practices for using HTTP Sessions	85

Managing HTTP sessions: Resources for learning:	88
Chapter 4. Using enterprise beans in applications	89
Enterprise beans.	89
Developing enterprise beans	90
Migrating enterprise bean code to the supported specification	91
WebSphere extensions to the Enterprise JavaBeans specification	93
Best practices for developing enterprise beans.	94
Unknown primary-key class.	98
Using access intent policies.	98
Access intent policies	98
Access intent service.	101
Access intent design considerations	102
Applying access intent policies to methods.	102
Using the AccessIntent API	103
Access intent exceptions	104
Access intent assembly settings	105
Access intent best practices	107
Frequently asked questions: Access intent	107
EJB modules	108
Assembling EJB modules	108
Container transactions	109
Method extensions	109
Method permissions	109
References	109
EJB containers	110
Container Managed Persistence over anything	110
Managing EJB containers	118
EJB container settings	118
EJB container system properties	119
EJB cache settings	119
Container interoperability	120
Deploying EJB modules	125
EJB module collection	125
EJB module settings	125
Enterprise beans: Resources for learning	126
EJB method Invocation Queuing	127
Chapter 5. Using extended messaging in applications	129
Extended messaging - overview	129
Extended messaging - receiving messages	130
Extended messaging - sending messages	131
Extended messaging - data mapping	132
Extended messaging - handling late responses	133
Extended messaging - transactional support	134
Extended messaging - exception handling	135
Extended messaging - application usage scenarios	136
Extended messaging - components	136
Designing an enterprise application to use extended messaging	138
Developing an enterprise application to use extended messaging	139
Deploying an enterprise application to use extended messaging	140
Configuring deployment attributes for a receiver bean.	140
Configuring deployment attributes for a sender bean	141
Configuring extended messaging service resources	143
Adding a new input port	143
Adding a new output port	143

Configuring an input port	144
Configuring an output port	144
Extended messaging service settings.	144
Extended messaging provider settings	147
Troubleshooting extended messaging	156
Extended Messaging: Resources for learning	156
Chapter 6. Using asynchronous messaging	159
Asynchronous messaging with WebSphere - an overview	159
Styles of messaging in applications	160
WebSphere Application Server cloning and WebSphere MQ clustering	160
Using JMS and messaging in applications	162
An overview of WebSphere asynchronous messaging using JMS	163
Administering WebSphere JMS support	166
Designing an enterprise application to use JMS	256
Developing a J2EE application to use JMS	259
Developing a JMS client	262
Deploying a J2EE application to use JMS	265
Tuning Java messaging service	266
Troubleshooting WebSphere Messaging.	268
Using message-driven beans in applications	272
Message-driven beans - an overview	272
Designing an enterprise application to use message-driven beans	275
Developing an enterprise application to use message-driven beans	276
Deploying an enterprise application to use message-driven beans	279
Configuring message listener resources for message-driven beans.	282
Important files for message-driven beans and extended messaging	291
Troubleshooting message-driven beans	292
Message-driven beans samples.	292
Chapter 7. Using application clients	295
Application clients	295
Application client functions.	296
ActiveX application clients	298
Applet clients	298
J2EE application clients.	299
Pluggable application clients	300
Thin application clients	302
Migration tips for application clients	303
Installing application clients	304
Best practices for installing application clients	306
Installing application clients silently	307
Developing application clients	308
Developing ActiveX application client code	308
Starting an ActiveX application	309
JClassProxy and JObjectProxy classes	311
Java virtual machine initialization tips.	314
Example: Developing an ActiveX application client to Enterprise Bean bridge using Java proxy objects	314
Example: Calling Java methods in the ActiveX to enterprise bean bridge.	315
Java field programming tips	316
ActiveX to Java primitive data type conversion values	317
Array tips for ActiveX application clients	318
Error handling codes for ActiveX application clients	319
Threading tips	319
Example: Viewing a System.out message	320

Example: Enabling logging and tracing for application clients	321
ActiveX client programming best practices	321
Developing applet client code	324
Accessing secure resources using SSL and applet clients	325
Applet client tag requirements	325
Applet client code requirements	326
Developing J2EE application client code	326
J2EE application client class loading	328
Developing pluggable application client code	330
Developing thin application client code	331
Developing thin application client code on a client machine	331
Developing thin application client code on a server machine	332
Assembling application clients	333
Deploying application clients on workstation platforms	334
Starting the Application Client Resource Configuration Tool and opening an EAR file	334
Data sources for application clients	334
Configuring new data source providers (JDBC providers) for application clients	335
Configuring new data sources for application clients	338
Configuring mail providers and sessions for application clients	338
Configuring new mail sessions for application clients	341
URLs for application clients	341
URL providers for the Application Client Resource Configuration Tool	342
Configuring new URL providers for application clients	342
Configuring new URLs with the Application Client Resource Configuration Tool	344
WebSphere asynchronous messaging using the Java Message Service API for the Application Client Resource Configuration Tool	345
Configuring Java messaging client resources	345
Configuring new connection factories for application clients	376
Configuring new Java Message Service destinations for application clients	376
Example: Configuring MQ Queue and Topic connection factories and destination factories for application clients	377
Example: Configuring WAS Queue and Topic connection factories and destination factories for application clients	378
Configuring new resource environment providers for application clients	379
Configuring new resource environment entries for application clients	380
Managing application clients	381
Updating data source and data source provider configurations with the Application Client Resource Configuration Tool	381
Updating URLs and URL provider configurations for application clients	382
Updating mail session configurations for application clients	382
Updating Java Message Service provider, connection factories, and destination configurations for application clients	382
Updating MQ Java Message Service provider, MQ connection factories, and MQ destination configurations for application clients	383
Updating resource environment entry and resource environment provider configurations for application clients	383
Removing application client resources	384
Running application clients	385
launchClient tool	386
Example: Using a Java 2 security manager with a J2EE application client	388
Example: Enabling Java 2 security prior to J2EE application client run time initialization	389
Application client troubleshooting tips	390
Chapter 8. Using Web services based on Web Services for J2EE	395
Web services	396
Web Services for J2EE	397

Java API for XML-based remote procedure call (JAX-RPC)	398
SOAP	398
SOAP with Attachments API for Java	399
Planning to use Web services based on Web Services for J2EE	399
Service-oriented architecture	400
Web services approach to a service-oriented architecture	400
Web services business models supported	403
Migrating Apache SOAP Web services to Web Services for J2EE	403
Developing Web services based on Web Services for J2EE	405
Example: Developing Web services based on Web Services for J2EE	406
Artifacts used to develop Web services based on Web Services for J2EE	408
Mapping between Java language, WSDL and XML	409
Java2WSDL command	429
WSDL2Java command	432
Setting up a development and unmanaged client execution environment for Web services based on Web Services for J2EE	435
Using the Java Messaging Service to transport Web services requests	435
Developing a Web service from a Java bean	437
Developing a Web service using a stateless session enterprise bean	448
Developing a new Web service with an existing WSDL file using a Java bean	451
Developing a new Web service from an existing WSDL file using a stateless session enterprise bean	452
Developing Web services clients based on Web Services for J2EE	454
Example: Developing Web services clients based on Web Services for J2EE	454
Developing client bindings from a WSDL file	455
Assembling a Web services-enabled client JAR file into an EAR file	456
Assembling a Web services-enabled client WAR file into an EAR file	457
Configuring the ibm-webservicesclient-bnd.xmi deployment descriptor	458
Configuring the webservicesclient.xml deployment descriptor	461
Configuring the webservicesclient.xml deployment descriptor for Handler classes	462
Testing Web services-enabled clients	466
Assembling Web services applications based on Web Services for J2EE	466
Assembling a Web services-enabled EJB JAR file	467
Assembling a Web services-enabled WAR file	469
Assembling a Web services-enabled EJB JAR into an EAR file	471
Assembling a Web services-enabled WAR into an EAR file	471
Enabling a Web services-enabled EAR file	472
Configuring the webservices.xml deployment descriptor	477
Configuring the ibm-webservices-bnd.xmi deployment descriptor	478
ibm-webservices-bnd.xmi assembly properties	479
Configuring the webservices.xml deployment descriptor for Handler classes	480
Deploying Web services based on Web Services for J2EE	481
wsdeploy command	481
Securing Web services based on WS-Security	483
Web services security specification- a chronology	484
Web services security support	485
Web services security and Java 2 Platform, Enterprise Edition security relationship	488
Web services security model in WebSphere Application Server	490
Web services security property collection	493
Web services security property configuration settings	493
Usage scenario for propagating security tokens	494
Configurations	495
Authentication method overview	509
XML digital signature	512
Securing Web services using XML digital signature	516
XML encryption	574

Securing Web services using XML encryption	577
Securing Web services using basicauth authentication	595
Identity assertion	602
Securing Web services using identity assertion authentication.	603
Securing Web services using signature authentication	608
Token type overview	613
Security token	619
Securing Web services using a pluggable token.	619
Tuning Web services based on Web Services for J2EE	629
Troubleshooting Web services based on Web Services for J2EE	629
Troubleshooting command-line tools for Web services based on Web Services for J2EE	630
Troubleshooting compiled bindings for Web services based on Web Services for J2EE	631
Troubleshooting the run time of Web services based on Web Services for J2EE	631
Troubleshooting the run time for a Web services client based on Web Services for J2EE	633
Troubleshooting serialization and deserialization in Web services based on Web Services for J2EE	633
Frequently asked questions about Web services based on Web Services for J2EE	635
Web services: Resources for learning	636
Chapter 9. Overview: Online garden retailer Web services scenarios	641
Web services scenario: Static inquiry on supplier	644
Web services scenario: Dynamic inquiry on supplier	646
Web services scenario: Cross supplier inquiry	648
Chapter 10. Web Services Invocation Framework (WSIF): Enabling Web services	651
Goals of WSIF	651
WSIF - Web services are more than just SOAP services	651
WSIF - Tying client code to a particular protocol implementation is restricting	652
WSIF - Incorporating new bindings into client code is hard	652
WSIF - Multiple bindings can be used in flexible ways	652
WSIF - Enabling a freer Web services environment promotes intermediaries	652
An overview of WSIF.	653
WSIF architecture	653
Using WSIF with Web services that offer multiple bindings	654
WSIF and WSDL	654
WSIF usage scenarios	654
Dynamic invocation	655
Using WSIF to invoke Web services	655
Using the WSIF providers	656
Developing a WSIF service	666
Using complex types.	675
Using the Java Naming and Directory Interface (JNDI)	676
Passing SOAP messages with attachments using WSIF.	678
Interacting with the J2EE container in WebSphere Application Server	680
Running WSIF as a client	680
WSIF system management and administration	681
Maintaining the WSIF properties file	681
Enabling security for WSIF	681
Troubleshooting the Web Services Invocation Framework	682
WSIF API	686
WSIF API reference: Creating a message for sending to a port	686
WSIF API reference: Finding a port factory or service.	687
WSIF API reference: Using ports	688
Chapter 11. Enabling the IBM WebSphere UDDI Registry	693
UDDI Registry terminology	693
UDDI Registry definitions	694

An overview of IBM UDDI Registries	694
Migrating from a previous version of IBM WebSphere UDDI Registry	695
Installing and setting up a UDDI Registry	698
Installing the UDDI Registry into a deployment manager cell	700
Installing the UDDI Registry into a single appserver	707
Reinstalling the UDDI Registry application	712
Applying Service updates to the UDDI Registry in a Network Deployment and single application server environment	713
Using a remote DB2 database for UDDI.	713
Removing the UDDI Registry application from a deployment manager cell	715
Removing the UDDI Registry application from a single application server	715
Configuring the UDDI Registry	716
Configuring global UDDI properties	716
Modifying the database userid and password	718
Configuring security roles	718
Configuring the UDDI User Console (GUI) for multiple language encoding support	719
Customizing the UDDI User Console (GUI)	719
Configuring SOAP interface properties	719
Configuring SOAP properties with the WebSphere Assembly Toolkit	720
Configuring SOAP properties in an application that is already deployed	720
Administering the UDDI Registry	720
Running the UDDI Registry	720
Backing up and restoring the UDDI Registry database	721
UDDI4J programming interface	721
UDDI user console	722
Displaying the user console	725
Custom Taxonomy Support in the UDDI Registry	726
UDDI Utility Tools	734
UDDI TestEntityExporter.java	745
UDDI TestEntityImporter.java	748
UDDI TestEntityPromoter.java	752
UDDI TestEntityFinder.java	754
UDDI TestEntityDeleter.java	758
UDDI TestUddiSerializer.java	761
UDDI TestUddiDeserializer.java	763
UDDI TestStubManager.java	765
UDDI TestCreateMinimalEntity.java	766
SOAP application programming interface for the UDDI Registry	768
Programming the SOAP API	768
SOAP API error handling tips in the UDDI Registry	769
UDDI Registry Application Programming Interface	769
Inquiry API for the UDDI Registry	769
Publish API for the UDDI Registry	771
UDDI EJB Interface for the UDDI Registry	772
Datatypes package in the UDDI Registry	777
EJB interface methods in the UDDI Registry	780
UDDI troubleshooting tips	781
Turning on UDDI trace	784
Messages	784
UDAI (Web Services UDDI) messages	785
UDCF (Web Services UDDI) messages	785
UDDA (Web Services UDDI) messages	786
UDDM (Web Services UDDI) messages	786
UDEJ (Web Services UDDI) messages	786
UDEX (Web Services UDDI) messages	786
UDIN (Web Services UDDI) messages	786

UDLC (Web Services UDDI) messages	805
UDPR (Web Services UDDI) messages	805
UDRS (Web Services UDDI) messages	805
UDSC (Web Services UDDI) messages	805
UDSP (Web Services UDDI) messages	805
UDUC (Web Services UDDI) messages	806
UDUT UDDI Utility Tools messages	808
UDUU (Web Services UDDI) messages	819
Running the UDDI samples	819
Installation Verification Program (IVP)	819
Reporting problems with the IBM WebSphere UDDI Registry	821
Feedback	821
Chapter 12. Web services gateway: Enabling Web services	823
Web services gateway - frequently asked questions	824
Web services gateway - What is new in this release	825
Web services gateway - Completing the installation	826
Web services gateway - prerequisites and constraints	827
Preserving an existing gateway configuration	828
Installing the gateway into an application server that is part of a deployment manager cell	829
Installing the gateway into a stand-alone application server	830
Testing the Web services gateway installation	832
Backing up and restoring a gateway configuration	833
Backing up and restoring UDDI publication links.	833
Backing up a gateway configuration	834
Restoring a gateway configuration	834
Administering the Web services gateway	835
Setting the namespace URI and WSDL URI for the Web services gateway	837
Working with channels	839
Working with filters	845
Working with JAX-RPC handlers	847
Working with UDDI references	851
Working with Web services	855
Running the Web services gateway Samples	866
Passing SOAP messages with attachments through the Web services gateway	866
SOAP messages with attachments - a definition.	867
Writing the WSDL extensions for SOAP messages with attachments	867
Developing Web services gateway extensions	868
Writing a filter for the Web services gateway	868
Using a filter to select a target service and port	877
Capturing Web service invocation information from the Web services gateway	878
Handling exceptions for the Web services gateway	879
Scaling the Web services gateway.	880
Scaling the Web services gateway using unlocked configuration deployment	881
Securing the Web services gateway	882
Enabling Web Services Security (WS-Security) for the gateway	882
Enabling basic authentication and authorization for the gateway	889
Invoking Web services over HTTPS	894
Enabling proxy authentication for the gateway	894
Tuning the Web services gateway	897
Selective SOAP parsing	899
Troubleshooting the Web services gateway	899
Web services gateway messages	904
Chapter 13. Class loading	923
Class loaders	923

Class loader collection	927
ClassLoader ID	927
ClassLoader Mode	927
Class loader settings	927
Migrating the class-loader Module Visibility Mode setting	927
Class loading: Resources for learning	928
Chapter 14. Using EJB query	929
EJB query language	929
Example: EJB queries	930
FROM clause	932
Inheritance in EJB query	933
Path expressions	933
WHERE clause	934
Scalar functions	941
Aggregation functions	944
SELECT clause	945
ORDER BY clause	946
Subqueries	946
EJB query restrictions	947
EJB Query: Reserved words	948
EJB query: BNF syntax	948
Comparison of EJB 2.0 specification and WebSphere query language	950
Using the dynamic query service	950
Example: Dynamic query remote client	952
Example: Dynamic query from local client	953
Dynamic query service performance considerations	953
Chapter 15. Using the internationalization service	955
Internationalization	955
Internationalization service: Overview	956
Internationalization challenges in distributed applications	957
Migrating internationalized applications	958
Assembling internationalized applications	959
Setting the internationalization type for servlets	959
Configuring container internationalization for servlets	960
Setting internationalization type for enterprise beans	961
Configuring container internationalization for enterprise beans	961
Using the internationalization context API	963
Gaining access to the internationalization context API	963
Accessing caller locales and time zone	964
Accessing invocation locales and time zone	965
Example: Internationalization context in an EJB client program	966
Example: Internationalization context in an EJB servlet	968
Example: Internationalization context in a session bean	970
Internationalization context API: Programming reference	971
Managing the internationalization service	980
Enabling the internationalization service for servlets and enterprise beans	981
Enabling the internationalization service for EJB clients	982
Troubleshooting the internationalization service	982
Internationalization service errors	983
Internationalization service exceptions	985
Internationalization: Resources for learning	986
Chapter 16. Application profiling	987
Application profiling: Overview	987

Application profiles	987
Application profiling performance considerations.	988
Application profiling tasks	990
Assembling applications for application profiling	991
Automatic configuration of application profiling	991
Applying profile-scoped access intent policies to entity beans	992
Creating a custom access intent policy	993
Creating an application profile	994
Configuring container managed tasks for application clients	995
Configuring container managed tasks for web components.	996
Configuring container managed tasks for Enterprise Java Beans	996
Configuring application managed tasks for application clients	997
Configuring application-managed tasks for web components	998
Configuring application managed tasks for Enterprise Java Beans	999
Specifying target servers for J2EE projects	999
Managing application profiles	1000
Using the TaskNameManager interface	1001
Application profiling exceptions	1002
Application profiling service settings.	1002
Application profile collection.	1003
Chapter 17. Using Business Rule Beans	1005
Advantages of externalizing business rules	1005
Overview of Business Rule Beans	1006
Externalized business rules	1007
Types of business rules	1007
Rule folders	1008
Rule attributes.	1008
Rule states	1010
Rule results.	1010
Dependent rules	1010
BRBeans run-time environment	1011
BRBeans run-time behavior	1011
BRBeans run-time exception handling	1012
Rule implementors	1012
Trigger point framework	1014
Trigger points	1015
As Of Date	1018
Predefined strategy objects	1018
Customized strategy objects	1020
Customized rule implementors.	1021
Rule management command	1022
Rule importer command	1022
Rule exporter command	1023
BRBeans properties file	1024
Database considerations for BRBeans	1025
Rule Management Application	1027
Rule management APIs	1027
BRBeans performance enhancements	1028
Developing BRBeans	1031
Determining where to place a trigger point	1031
Placing a trigger point in the application code	1032
Administering strategy objects to control triggers	1033
Implementing business rules	1034
Assembling applications for use with BRBeans.	1035
Managing rules	1036

Starting the Rule Management Application	1037
Copying or moving rules or rule folders	1037
Working with Quick Copy.	1037
Finding a rule	1038
Deleting rules	1038
Deleting rule folders	1038
Changing the properties of a rule	1039
Importing a rule	1039
Exporting a rule	1039
Renaming rules	1040
Renaming rule folders	1040
Specifying columns	1040
Changing the date and time format	1040
Rule Browser	1041
File menu	1041
Edit menu	1054
View menu	1057
Find Rules window	1058
Business rule beans: Resources for learning	1062
Chapter 18. Using asynchronous beans	1065
Asynchronous beans	1065
Example: Asynchronous bean connection management	1067
Configuring work managers	1068
Work managers	1069
Work manager collection	1070
Work manager service settings	1072
Assembling applications that use work managers	1073
Developing work objects to run code in parallel	1073
Work objects	1074
Example: Work object	1075
Developing event listeners	1076
Using the application notification service	1076
Example: Event listener	1077
Developing asynchronous scopes	1078
Asynchronous scopes	1079
Alarms	1080
Subsystem monitors	1081
Asynchronous scopes: Dynamic message bean scenario	1081
Interoperating with asynchronous beans	1082
Serialized Asynchronous beans interoperability.	1082
Java Management Extension MBean interoperability.	1084
Chapter 19. Using object pools	1085
Object pool managers	1086
Object pool manager collection	1087
Name	1087
JNDI Name.	1087
Description	1088
Category.	1088
Object pool manager settings	1088
Object pool service settings	1089
Startup	1089
Object pools: Resources for learning	1089
Object pool performance considerations	1090

Chapter 20. Using startup beans	1093
Chapter 21. Using the scheduler service	1095
Managing the scheduler service	1095
Creating the database for a scheduler	1095
Configuring a scheduler	1100
Enabling the scheduler service.	1104
Developing and scheduling tasks	1105
Developing a task that calls a session bean	1106
Developing a task that sends a JMS message	1106
Receiving scheduler notifications	1107
Submitting a task to a scheduler	1108
Managing tasks with a scheduler	1109
Scheduler interface	1110
Interoperating with the Scheduler service	1113
Recreating Scheduler tasks	1113
Deleting Scheduler tasks	1115
Recreating Scheduler tables.	1115
Chapter 22. Using shared work areas	1117
WorkArea service - Overview	1117
Work area property modes	1117
Nested work areas	1118
Distributed work areas.	1120
WorkArea service: Special considerations.	1121
Work area service performance considerations.	1121
Developing applications that use work areas	1122
UserWorkArea interface	1123
Example: WorkArea SimpleSample application.	1123
Accessing the WorkArea service	1124
Beginning a new work area	1124
Setting properties in a work area	1125
Using a work area to manage local work	1126
Completing a work area	1129
Managing the work area service	1131
Enabling the WorkArea service	1131
Managing the size of work areas	1132
Configuring work area partitions on the server	1133
Work area partition service	1133
WorkAreaPartitionManager interface.	1135
Example: Work area partition manager.	1136
Work area partition collection	1137
Bidirectional propagation	1139
Accessing a user defined work area partition	1140
Chapter 23. Working with the Event Programming Model in WebSphere.	1143
The Common Event Infrastructure in WebSphere	1143
Events and Common Base Events	1144
The structure of the Common Base Event	1144
Components you should know about	1149
Starting to use the Event Programming Model	1150
Planning to use the Event Programming Model.	1151
Installing the Event Programming Model	1153
Configuring the Event Programming Model using the Administration console	1154
Generating Common Base Events	1158
Obtain an emitter.	1159

Event creation from WebSphere applications	1159
Using the Java API	1159
Generating a Common Base Event in BPEL	1160
Populating events with event context data	1161
Send an event.	1163
Viewing events with the event browser	1163
Specifying the events to view	1164
Working with the events returned	1166
Working with WebSphere event infrastructure messages	1168
State observer plugin messages	1168
Business context data services messages	1171
Common Event Infrastructure Developer's Guide	1173
Introduction to the common event infrastructure	1173
Configuring the common event infrastructure	1175
Developing an event source.	1180
Developing an event consumer	1184
Working with events	1188
Message reference	1193
CEIC0001W	1193
CEIC0002W	1193
CEIC0003W	1193
CEIC0004W	1193
CEIC0005W	1194
CEIC0006I	1194
CEIDS0001E	1194
CEIDS0002E	1194
CEIDS0003E	1195
CEIDS0004I	1195
CEIDS0005E	1195
CEIDS0006E	1195
CEIDS0007E	1196
CEIDS0008E	1196
CEIDS0009E	1196
CEIDS0010I	1196
CEIDS0011E	1197
CEIDS0013E	1197
CEIDS0014E	1198
CEIDS0015E	1198
CEIDS0016E	1199
CEIDS0017E	1199
CEIDS0018E	1199
CEIDS0019E	1200
CEIDS0020E	1200
CEIDS0021E	1200
CEIDS0022E	1201
CEIDS0023E	1201
CEIDS0024E	1201
CEIDS0025E	1202
CEIDS0026E	1202
CEIDS0027E	1202
CEIDS0028E	1203
CEIDS0029E	1203
CEIDS0030E	1203
CEIDS0031I	1203
CEIDS0032I	1204
CEIDS0033I	1204

CEIDS0034E	1204
CEIEI0001E	1204
CEIEI0002E	1205
CEIEI0003E	1205
CEIEI0005E	1205
CEIEI0007E	1206
CEIEI0008E	1206
CEIEI0010E	1206
CEIEI0011E	1207
CEIEI0015E	1207
CEIEI0016E	1207
CEIEI0017E	1208
CEIEI0018E	1208
CEIEI0019E	1208
CEIEI0020E	1209
CEIEI0021E	1209
CEIEI0022E	1209
CEIEI0023E:	1210
CEIEI0024E	1210
CEIEI0025E	1210
CEIEI0026E	1211
CEIEI0027E	1211
CEIEI0028E	1211
CEIEI0029E	1212
CEIEI0030E	1212
CEIEI0031E	1212
CEIEM0001E	1213
CEIEM0002E	1213
CEIEM0003E	1213
CEIEM0004E	1213
CEIEM0005E	1214
CEIEM0006E	1214
CEIEM0007E	1214
CEIEM0008E	1214
CEIEM0014E	1215
CEIEM0015E	1215
CEIEM0020E:	1215
CEIEM0021E	1215
CEIEM0022E	1216
CEIEM0023E	1216
CEIEM0024E	1216
CEIEM0025E	1216
CEIEM0030E	1217
CEIEM0031E	1217
CEIEM0032E	1217
CEIEM0033E	1217
CEIEM0034E	1218
CEIEM0035E	1218
CEIEM0036E	1218
CEIES0001E	1218
CEIES0002E	1219
CEIES0003E	1219
CEIES0004E	1219
CEIES0005E	1220
CEIES0006E	1220
CEIES0007W	1220

CEIES0008W	1220
CEIES0009E	1221
CEIES0010E	1221
CEIES0011E	1221
CEIES0012E	1221
CEIES0013E	1222
CEIES0014W	1222
CEIES0015E	1222
CEIES0016E	1222
CEIES0018E	1223
CEIES0019E	1223
CEIES0020E	1223
CEIES0021E	1223
CEIES0023E	1224
CEIES0024E	1224
CEIES0025E	1224
CEIES0026E	1224
CEIES0027E	1225
CEIES0028E	1225
CEIES0029E	1225
CEIES0030W	1226
CEIES0031E	1226
CEIES0032E	1226
CEIES0033E	1226
CEIES0034E	1227
CEIES0035E	1227
CEIES0036E:	1227
CEIES0037E	1227
CEIES0038E	1228
CEIES0039E	1228
CEIES0040E	1228
CEIES0041E	1228
CEIES0042E	1229
CEIES0043E	1229
CEIES0044E	1229
CEIES0045E	1229
CEIES0046E:	1230
CEIES0048E	1230
CEIES0049E	1230
CEIES0050E	1230
CEIIN0001E	1231
CEIIN0002I	1231
CEIIN0003E	1231
CEIIN0004I	1231
CEIIN0005E	1231
CEIIN0006E	1232
CEIIN0007E	1232
CEIIN0008E	1232
CEIIN0009E	1232
CEIIN0010E	1233
CEIIN0011E	1233
CEIIN0012E	1233
CEIIN0013E	1233
CEIIN0014E	1233
CEIIN0015E	1234
CEIIN0016E	1234

CEIIN0017E	1234
CEIIN0018I	1234
CEIIN0019W	1235
CEIIN0020E	1235
CEIIN0021E	1235
CEIIN0022E	1235
CEIIN0023E	1235
CEIIN0024E	1236
CEIIN0025E	1236
CEIIN0026I	1236
CEIIN0027E	1236
CEIIN0028E	1237
CEIIN0029E	1237
CEIIN0030E	1237
CEIIN0031E	1237
CEIIN0032E	1237
CEIIN0033E	1238
CEIIN0034E	1238
CEIIN0035W	1238
CEIIN0036W	1238
CEIIN0037E	1238
CEIIN0038E	1239
CEIIN0039E	1239
CEIIN0040E	1239
CEIIN0041E	1239
CEIIN0042I	1240
CEIIN0043E	1240
CEIIN0044E	1240
CEIIN0045I	1240
CEIIN0046E	1240
CEIIN0047E	1241
CEIIN0048E	1241
CEIIN0049I	1241
CEIIN0050W	1241
CEIIN0051I	1241
CEIIN0052I	1242
CEIIN0053I	1242
CEIIN0054E	1242
CEIIN0055I	1242
CEIIN0056E	1242
CEIIN0057I	1242
CEIIN0058I	1243
CEIIN0059I	1243
CEIIN0060E	1243
CEIIN0061W	1243
CEIIN0062I	1243
CEIIN0063E	1244
CEIIN0064E	1244
CEIIN0065I	1244
CEIIN0066E	1244
CEIIN0067E	1244
CEIIN0068E	1245
CEIIN0069E	1245
CEIIN0070I	1245
CEIIN0071I	1245
CEIIN0072E	1245

CEIIN0073I	1246
CEIIN0074E	1246
CEIIN0075I	1246
CEIIN0076E	1246
CEIIN0077E	1246
CEIIN0078E	1247
CEIIN0079E	1247
CEIIN0083E	1247
CEIIN0084E	1247
CEIIN0085I	1248
CEIIN0086E	1248
CEIIN0087I	1248
CEIIN0088E	1248
CEIIN0089E	1249
CEIIN0090E	1249
CEIIN0091E	1249
CEIIN0092I	1249
CEIIN0093I	1250
CEIIN0094I	1250
CEIIN0095E	1250
CEIIN0096E	1250
CEIIN0097E	1250
CEIIN0098E	1251
CEIIN0099I	1251
CEIIN0100E	1251
CEIIN0101I	1251
CEIIN0102I	1252
CEIIN0103I	1252
CEIIN0104E	1252
Chapter 24. Using the transaction service	1253
Transaction support in WebSphere Application Server	1253
Resource manager local transaction (RMLT)	1254
Global transactions	1254
Local transaction containment (LTC)	1255
Local and global transaction considerations	1256
Extended JTA support	1256
Developing components to use transactions.	1257
Configuring transactional deployment attributes using the Assembly Toolkit	1257
Using bean-managed transactions	1259
Configuring transaction properties for an application server	1260
Transaction service settings.	1261
Using local transactions	1263
Managing active transactions	1265
Managing transaction logging for optimum server availability.	1266
Configuring transaction aspects of servers for optimum availability	1267
Moving a transaction log from one server to another.	1268
Restarting an application server on a different host	1269
Interoperating transactionally between application servers.	1269
Troubleshooting transactions	1269
Transaction service exceptions	1270
UserTransaction interface - methods available	1271
Using one-phase and two-phase commit resources in the same transaction	1271
Coordinating access to 1-PC and 2-PC-capable resources within the same transaction	1271
Assembling an application to use one-phase and two-phase commit resources in the same transaction	1272

Configuring an application server to log heuristic reporting	1274
Exceptions thrown for transactions involving both single- and two-phase commit resources	1274
Last Participant Support: Resources for learning	1274
Chapter 25. Using the ActivitySession service	1277
The ActivitySession service	1277
Using ActivitySessions with HTTP sessions	1278
ActivitySession and transaction contexts	1280
Combining transaction and ActivitySession container policies	1281
Developing a J2EE application to use ActivitySessions	1286
Developing an enterprise bean or J2EE client to manage ActivitySessions	1287
Setting EJB module ActivitySession deployment attributes with the Assembly Toolkit	1288
Setting Web module ActivitySession deployment attributes with the Assembly Toolkit	1291
Disabling or enabling the ActivitySession service	1292
Configuring the default ActivitySession timeout for an application server	1293
ActivitySession service settings	1294
Troubleshooting ActivitySessions	1294
The ActivitySession service application programming interfaces	1295
Samples: ActivitySessions	1295
ActivitySession service: Resources for learning	1297
Chapter 26. Using naming	1299
Naming	1299
Version 5 features for name space support	1300
Name space logical view	1301
Initial context support	1303
Lookup names support in deployment descriptors and thin clients	1304
JNDI support in WebSphere Application Server	1306
Developing applications that use JNDI	1306
Example: Getting the default initial context	1308
Example: Getting an initial context by setting the provider URL property	1311
Example: Setting the provider URL property to select a different root context as the initial context	1313
Example: Looking up an EJB home with JNDI	1314
Example: Looking up a JavaMail session with JNDI	1316
JNDI interoperability considerations	1316
JNDI caching	1317
JNDI cache settings	1318
Example: Controlling JNDI cache behavior from a program	1319
JNDI name syntax	1320
INS name syntax	1320
JNDI to CORBA name mapping considerations	1321
Example: Setting the syntax used to parse name strings	1321
Developing applications that use CosNaming (CORBA Naming interface)	1322
Example: Getting an initial context with CosNaming	1322
Example: Looking up an EJB home with CosNaming	1324
Configured name bindings	1326
Name space federation	1328
Name space bindings	1329
Configuring and viewing name space bindings	1329
String binding settings	1329
CORBA object binding settings	1330
Indirect lookup binding settings	1331
EJB binding settings	1331
Name space binding collection	1332
Configuring name servers	1333
Name server settings	1333

Troubleshooting name space problems	1333
dumpNameSpace tool	1334
Example: Invoking the name space dump utility	1336
Name space dump utility for java:, local: and server name spaces	1336
Example: Invoking the name space dump utility for java: and local: name spaces	1338
Name space dump sample output	1339
Naming and directories: Resources for learning	1341
Chapter 27. Using the dynamic cache service to improve performance	1343
Dynamic cache	1343
Configuring cache replication	1343
Cache replication	1344
Internal messaging configuration settings	1344
Enabling the dynamic cache service	1345
Dynamic cache service settings	1345
Configuring servlet caching	1346
Configuring the dynamic cache disk offload	1346
Configuring Edge Side Include caching	1347
Configuring external cache groups	1349
Displaying cache information	1353
Configuring cacheable objects with the cachespec.xml file	1354
Verifying the cacheable page	1355
Cachespec.xml file	1356
Configuring command caching	1362
Command class	1362
CacheableCommandImpl class	1362
Example: Caching a command object	1363
Using the DistributedMap interface for the dynamic cache	1364
Sharing cached objects in a clustered environment	1366
Cache instance settings	1366
Cache instance collection	1367
Cache instance service settings	1368
Invalidation listeners	1368
Example: Caching Web services	1369
Example: Configuring the dynamic cache	1371
Cache monitor	1373
Edge cache statistics	1374
Troubleshooting the dynamic cache service	1375
Troubleshooting tips for the dynamic cache service	1376
Chapter 28. Using user profiles.	1379
User profile	1379
UserProfileManager class	1379
User profile development options	1380
Extending the data represented in user profiles	1380
Adding columns to the base user profile implementation	1380
Extending the User Profile enterprise bean and importing legacy databases	1381
Example: UPServletExample.java	1381
Example: UserProfileExtendedSample.java	1382
Example: UPServletExampleExtended.java	1383
Example: UserProfileExtended.java	1385
Example: UPServletExtended.java	1386
userprofile.xml	1388
Chapter 29. Assembling applications with the Assembly Toolkit	1391
Application assembly and J2EE applications	1393

Archive support in Version 5	1393
Starting the Assembly Toolkit	1393
astk command	1394
Migrating code artifacts to the Assembly Toolkit	1395
Importing enterprise applications	1395
Importing WAR files.	1395
Importing client applications.	1396
Importing EJB files	1396
Importing RAR files or connectors	1397
Creating enterprise applications	1398
Creating Web applications	1398
Creating application clients	1400
Creating EJB modules.	1401
Creating connector modules	1402
Editing deployment descriptors	1403
Mapping enterprise beans to database tables	1404
Mapping constraints for databases	1404
Verifying archive files	1405
Generating code for EJB deployment	1406
Generating code for Web service deployment	1406
Assembly Toolkit: Resources for learning	1407
Chapter 30. Deploying and managing applications	1409
Enterprise applications	1409
Installing a new application	1409
Preparing for application installation settings	1413
Example: Installing an EAR file using the default bindings.	1417
Enterprise application collection	1417
Name	1417
Status.	1417
Enterprise application settings	1417
Starting and stopping applications	1421
Exporting applications	1421
Exporting DDL files	1422
Updating applications	1422
Hot deployment and dynamic reloading	1423
Uninstalling applications	1431
Deploying and managing applications: Resources for learning	1431
Notices	1433
Trademarks and service marks	1435

How to send your comments

Your feedback is important in helping to provide the most accurate and highest quality information.

- To send comments on articles in the WebSphere Application Server Information Center
 1. Display the article in your Web browser and scroll to the end of the article.
 2. Click on the **Feedback** link at the bottom of the article, and a separate window containing an e-mail form appears.
 3. Fill out the e-mail form as instructed, and click on **Submit feedback** .
- To send comments on PDF books, you can e-mail your comments to: **wasdoc@us.ibm.com** or fax them to 919-254-0206.

Be sure to include the document name and number, the WebSphere Application Server version you are using, and, if applicable, the specific page, table, or figure number on which you are commenting.

When you send information to IBM, you grant IBM a nonexclusive right to use or distribute the information in any way it believes appropriate without incurring any obligation to you.

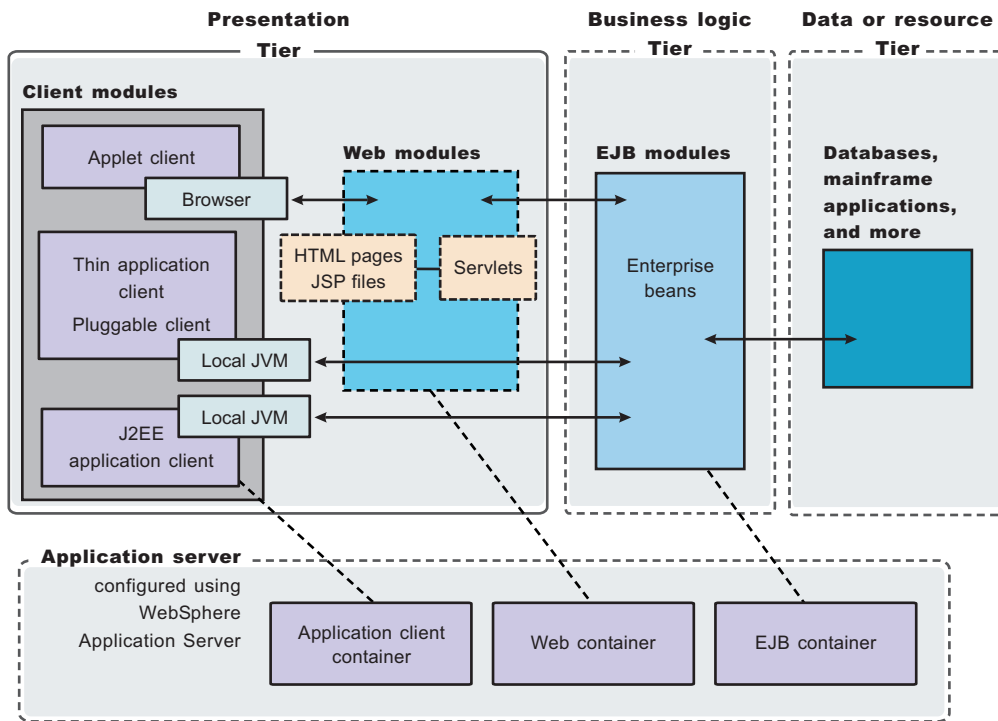
Chapter 1. Welcome to applications

In the WebSphere Application Server environment, the components making up a Java 2 Platform, Enterprise Edition (J2EE) application are referred to as *application modules*, specifically:

- Web modules
- Enterprise JavaBeans (EJB) modules
- Client modules

After you assemble modules into a main module called an application module or EAR file, you are ready to deploy the application module onto the application server.

You can deploy both Web and EJB modules as stand-alone applications. A J2EE application is generally comprised of all three application modules, each deployed in its own container that is provided by the application server that you have configured. The following diagram depicts how the application modules interact across the three-tier architecture of enterprise computing:



The diagram does not depict interaction between Web modules and the thin application, pluggable, or J2EE application client because it is a rare requirement. Applications that use these client types, however, can incorporate Web modules. The product provides different levels of support for the requirement, according to client type. Refer to articles specifically on your client type for more information.

Web modules

A Web module consists of the following components, in varying combinations (corresponding to the application client requirements, security requirements, and so on):

- HyperText Markup Language (HTML) pages and JavaServer Pages (JSP) supply the user interface and program logic.
- Servlets coordinate work between these components and business logic.

You create a Web module by assembling servlets, JSP files, and static content such as HTML pages into a single deployable unit. You can then deploy the Web module as a stand-alone Web application, or combine it with other modules to create J2EE applications.

Whichever application type you deploy, you have the option of activating **Session Manager** for your Web module. By using this service to manage HTTP sessions, you can personalize Web pages for individual users. A session is a series of requests to a servlet, originating from the same user at the same browser. Managing HTTP sessions allows servlets running in a Web container to keep track of individual users.

For example, a servlet might use sessions to provide "shopping carts" to on-line shoppers. Suppose the servlet is designed to record the items each shopper indicates he or she will purchase from the Web site. It is important that the servlet be able to associate incoming requests with particular shoppers. Otherwise, the servlet might mistakenly add choices of Shopper 1 to the cart of Shopper 2.

EJB modules

EJB modules are combinations of enterprise beans, which perform various roles in retrieving data from the Data and Resource tier and running business logic in a J2EE application. You also can deploy an EJB module as a stand-alone application (refer to Chapter 30, "Deploying and managing applications," on page 1409 for the necessary steps).

Version support

IBM WebSphere Application Server provides broad support for enterprise beans, including the Enterprise JavaBeans (EJB) 2.0 specification. The EJB 2.0 specification introduces a container-managed persistence (CMP) 2.0 component model, which provides a number of improvements to aid developer productivity and application performance. In addition, this product continues to fully support enterprise beans written to the CMP 1.1 programming model and deployed in previous versions of this product; applications can use CMP 1.1 beans, CMP 2.0 beans, or a mixture of both. CMP 1.1 beans can be directly carried forward in an EJB 1.1 ejb-jar module or may be repackaged and combined with CMP 2.0 beans in an EJB 2.0 module.

Several excellent trade books that cover the latest EJB specification and the CMP persistence model are already available. A good way to locate some of these is to visit your favorite online bookstore and search on the term *Enterprise JavaBeans*. For a more basic orientation, see "Enterprise beans: Resources for learning" on page 126.

Innovations for EJB development

For EJB 2.0 modules, a feature introduced in Version 5 of this product, called **access intent** policies, eases the management of interactions between CMP beans and their underlying data stores. Each policy sets such data access characteristics such as access type (read or update) and transaction isolation that affect the locking of resources, letting you choose the level of data integrity and performance for your application. The Integration Server product adds APIs to enable you to further customize IBM-provided access intent policies for your particular environment.

Your application development can also include **asynchronous messaging**, which the product supports as a method of communication based on the Java Message Service (JMS) programming interface. The base JMS support enables IBM WebSphere Application Server applications to exchange messages asynchronously with other JMS clients by using JMS destinations (queues or topics). An application can explicitly poll for messages on a destination.

The product also provides a message listener service that applications can use to automatically retrieve messages from JMS destinations for processing by message-driven beans, without the application having to explicitly poll JMS destinations.

The Integration Server product provides **extended messaging**, which uses the EJB container to manage the messaging infrastructure, and provides more types of messaging beans. This enables application developers to concentrate on the business logic for enterprise beans and to leave the messaging usage to messaging objects and configuration of the EJB container.

Refer to:

- “Asynchronous messaging with WebSphere - an overview” on page 159 An overview of WebSphere asynchronous messaging.
- “Using JMS and messaging in applications” on page 162 Implementing WebSphere J2EE applications that use JMS.

For information about the types of enterprise beans, see the Enterprise beans article.

Client modules

WebSphere Application Server supports multiple models of a client module, each suitable for different types of applications. For example: the various dependencies of the models on Web modules for access to business logic can accommodate different client-side requirements for deployment. If a developer wants to create an application requiring no initial code deployment to client machines, an appropriate client module choice is the applet client. This client receives the necessary application code through download from HTML pages.

Most models of a client module are only available when you install the product. For a list of supported products and version information, see the Supported Prerequisites Web site.

ActiveX application client

WebSphere Application Server provides an ActiveX to EJB bridge that enables ActiveX programs to access WebSphere Enterprise JavaBeans through a set of ActiveX automation objects.

Applet client

The applet client provides a browser-based Java run time capable of interacting with enterprise beans directly, instead of indirectly through a servlet.

J2EE application client

The J2EE application client programming model provides the benefits of Java TM 2 Platform for WebSphere Application Server Enterprise (J2EE).

Pluggable application client (Windows platforms only)

The pluggable application client is a thin application client that uses the Sun Java run-time environment (JRE) instead of the IBM JRE environment.

Thin application client

This application client is a Java application client without the J2EE client container support. It uses a different programming model than the J2EE application client.

5.1 +

Integration Server client

The Integration Server client is a thin application client that provides additional run time support for the J2EE application client and the thin application client to access WebSphere Business Integration Server Foundation server features, such as business rule beans, internationalization services, and the activity session service.

Note: The Imaginary Buffer Line product provides a CD-ROM and installation program for installing application clients without installing the entire application server. This smaller footprint is useful if you want to run client applications on multiple client machines. For more information, see Installing application clients.

Web services

Web services are services that you use over the Internet. If you have an existing application, and you want to make the service that your application provides available to others - either within your own organization or beyond it - you can use Web services technologies to provide a standard Web interface for your service. Used in this manner, Web services can be defined as middleware. You can connect applications together no matter how each application is implemented or where it is located.

Middleware is not new, but what is new is Web services technology and its power to connect by using open standards. Web services operate at a level of abstraction that is similar to the Internet; they can work with any operating system, hardware platform or programming language that can be Web-enabled.

The core technologies on which Web services are developed and implemented include:

- XML (Extensible Markup Language). XML solves the problem of data independence. You use it to describe data, and also to map that data into and out of any application or programming language.
- Web services for Java 2 platform, Enterprise Edition (J2EE) specification defines the programming model and run-time architecture for implementing Web services based on the Java language. WebSphere Application Server Version 5.0.2 and 5.1 supports Web Services for J2EE Version 1.0. If you want to know how to implement a Web service interface to an existing application, then deploy your Web service within the application server, see Using Web services based on Web Services for J2EE.
- Java API for XML-Based RPC (JAX-RPC) enables you to develop SOAP-based interoperable and portable Web services and Web services clients. WebSphere Application Server Version 5.0.2 and 5.1 supports JAX-RPC Version 1.0.
- WSDL (Web Services Description Language). You use this XML-based language to create a description of an underlying application. It is this description that turns an application into a Web service, by acting as the interface between the underlying application and other Web-enabled applications. WebSphere Application Server Version 5.0.2 and 5.1 supports WSDL Version 1.1.
- SOAP (Simple Object Access Protocol). SOAP is the core communications protocol for the Web, and most Web services use this protocol to talk to each other. WebSphere Application Server Version 5.0.2 and 5.1 supports SOAP Version 1.1.
- SOAP with attachments API for Java (SAAJ) is used for SOAP messaging that works behind the scenes in the Java API for XML-based RPC (JAX-RPC) implementation. You can also use this API to directly write SOAP messaging applications rather than using JAX-RPC. SAAJ allows you to do XML messaging from the Java platform by making method calls by creating, sending and consuming XML messages over the Internet. WebSphere Application Server Version 5.0.2 and 5.1 supports SAAJ Version 1.1.

WebSphere Application Server also provides other mechanisms that can help you get the most out of your Web services:

A Web Services Invocation Framework (WSIF)

SOAP bindings for Web services are part of the WSDL specification. So when you think of using a Web service, you probably think of assembling a SOAP message and sending it across the network to the service endpoint, using some SOAP client API. The WSDL specification allows for extensibility points which can describe alternate ways of invoking a Web service. A WSIF client can make use of these non-SOAP descriptions to invoke a service in a more efficient way. For example, a Web service provider might offer a SOAP binding for the service and a local Java binding that allows you to treat the local service implementation (a Java class) as a Web service. If the client is deployed in the same environment as the service, then the local Java binding for the service can be used. This provides more efficient communication with the service by making direct Java calls rather than sending and receiving SOAP messages.

To deploy a Web service as a WSIF-enabled service, you first develop and deploy the Web service, then you develop the WSIF client based on the WSDL document for that Web service -

unless you are also using the Web services gateway, in which case the gateway automatically redeploys your Web service as a WSIF-enabled service.

A Web services gateway

You use the gateway to handle Web service invocations between Internet and Intranet environments. You use it to make your internal Web services available externally, and to make external Web services available to your internal systems. You also use it to specify:

- The transport mechanisms (or channels) on which messages can be carried to and from the service.
- The filters or handlers (if any) that act upon these incoming and outgoing messages.
- The UDDI registries (if any) to which you want the service to be published
- The levels of security that you want to apply to the service.

When you deploy a Web service to the gateway, the gateway creates a copy of the WSDL file for that service and stores it at a new Web address. Users of the service through the gateway then use the gateway copy of the WSDL file. You should therefore (if possible) decide whether or not you want to use the gateway before you make the Web addresses of your deployed services available to others.

CORBA applications 5.1 +

In addition to J2EE enterprise applications, you can also deploy CORBA client/server applications in your environment (if your environment includes the CORBA C++ SDK). You can use the CORBA C++ SDK to build a lightweight WebSphere CORBA server to use with new or existing C and C++ programs. You also can use the SDK to build a WebSphere CORBA C++ client to use with a WebSphere EJB server or WebSphere CORBA C++ server.

For more information, refer to CORBA applications documentation.

Note: **5.1 +** In future releases, IBM intends to remove the C++ Object Request Broker (ORB), the C++ library for IDL valuetypes and the WebSphere Application Server C++ security client. IBM will no longer ship or support the Common Object Request Broker Architecture (CORBA) C++ Developer Kit. The CORBA technology is a bridge for migration to a Java 2 Platform Enterprise Edition (J2EE) and WebSphere Application Server environment. It is recommended that customers migrate to the Object Request Broker (ORB) service for Java technology that ships with WebSphere Application Server. However, there is no equivalent J2EE functionality for the C++ security client or the C++ Valuetype library. Customers that require such functionality must provide or develop their own.

The deprecation of the CORBA C++ Developer Kit does not affect support for CORBA interoperability with vendor software for CORBA services. View the following links for additional information about interoperability:

- [CORBA Interoperability Samples documentation](#)
- [IBM WebSphere Application Servers CORBA Interoperability white paper](#)

Application services

IBM WebSphere Application Server provides essential services to ease the building of dynamic and flexible e-business applications. These services support and extend the open standards of J2EE and Web services, with a focus on application reuse and integration.

The Enterprise Extensions product takes application services to the next level, providing a broad range of dynamic API extensions that address functional gaps in the J2EE programming model.

- **Class loading**

The WebSphere Application Server product provides several class-loading modes, policies, and features to enable you to deploy and run your applications successfully. An application server provides an Application class-loader policy that enables you to control the isolation of applications in a server. If you want applications to share classes, choose the SINGLE policy; otherwise choose the MULTIPLE policy, which isolates the class loaders for each application.

Similarly, at the application level, you can choose a WAR class-loader policy that configures the isolation of Web modules within an application. If you choose the policy APPLICATION, then each Web module in your application can see the classes of other Web modules. A policy of MODULE creates a separate class loader for each Web module, resulting in isolation for each of the classes of each Web module.

The class-loader mode setting, which you can configure at the server, application, or Web module level depending on your class-loader policy, enables you to control whether application class loaders override classes contained in base run-time class loaders. By default, the WebSphere Application Server class loaders have a class-loader mode of PARENT_FIRST, which is the standard JDK mode and does not allow the application class loader to override classes. You must take care when using the PARENT_LAST class-loader mode to make all dependent classes available within the application or you might get LinkageErrors or other class-loader exceptions. For example, if you provide a newer version of the Xerces.jar file and your application is using XSLT, you must also provide a xalan.jar file within your application.

- **Shared library**

Version 5.0 of WebSphere Application Server introduces the concept of a shared library. A shared library is a CLASSPATH and a symbolic name for that classpath. You define shared libraries at the cell, node, or server level and then associate the shared libraries either with an application server (making the classes available to all applications in the server) or with individual applications (making the classes available only to the referencing application). This mechanism provides a convenient way to make libraries of classes available to your applications outside of a standard J2EE enterprise application (EAR) file for easier version management and space efficiency. (See "Managing shared libraries" in the Information center.)

- **EJB query**

The EJB query language is used to specify a query over container-managed entity beans. WebSphere's EJB query language is compliant with the EJB QL defined in Sun's EJB 2.0 specification, but adds additional support as described in the topic "Comparison of EJB 2.0 specification and WebSphere query language" on page 950.

EJB query can be used to define a finder or select method of an EJB entity bean. Finder and select queries are specified in the bean deployment descriptor using the <ejb-ql> tag. Queries specified in the deployment descriptor are compiled into SQL during deployment. See Chapter 14, "Using EJB query," on page 929.

The Enterprise Extensions product includes the dynamic query service, an additional API that enables you to dynamically specify a query in your application by adding the executeQuery() method. See "Using the dynamic query service" on page 950.

- **Internationalization support**

The internationalization service manages the distribution of locale and time zone information, or *internationalization context*, in applications that run on WebSphere Application Server Enterprise installations. The internationalization service solves the problem of mismatched locales and time zones by systematically managing the distribution of internationalization context across the various components of EJB applications.

The internationalization service transparently propagates internationalization context over requests that originate from J2EE-compliant Web service clients. The service creates a SOAP header block that contains the invocation context scoped to the current thread; this SOAP representation is then inserted into the outgoing Web services request. For incoming requests, the service scopes the propagated internationalization context, referred to as *caller context*, to the invocation of the stateless session bean that is enabled as a Web service. The service also scopes an *invocation context* as prescribed by the internationalization context management policies that were assigned to the enterprise bean's methods during application assembly.

For more information about the internationalization service, see Chapter 15, “Using the internationalization service,” on page 955

- **Workarea service**

The WorkArea service enables application developers to implicitly propagate information beyond the information passed in remote calls. Applications can create a work area, insert information into it, and make remote invocations. The work area is propagated with each remote method invocation, eliminating the need to explicitly include an appropriate argument in the definition of each method. The methods on the server side can use or ignore the information in the work area as appropriate. See Chapter 22, “Using shared work areas,” on page 1117

- **Application profiling**

Application profiling enables you to configure multiple access intent policies on the same method of an entity bean; and to configure multiple access intent policies for dynamic query on the same entity bean.

To use application profiling, application developers identify named units of work, or *tasks*. A task typically corresponds to the execution of a concrete and high-level job within the application. The IBM WebSphere Application Server run-time environment queries the task at the invocation of any entity bean, and establishes the appropriate access intent policy under which the bean should execute. An *application profile* is the set of access intent or query intent policies that should be selectively applied, as well as the list of tasks for which the policies should be applied. See Chapter 16, “Application profiling,” on page 987

- **Scheduler service**

The scheduler service enables J2EE work to be executed at a requested time or interval. The scheduler API supports different implementations of the TaskInfo interface, each of which can be used to schedule a particular type of work; for example, you can develop a task that calls a session bean or a task that sends a JMS message. You can set a notification sink on a task in order to receive the notification events that are generated by a scheduler when it performs an operation on the task. See Chapter 21, “Using the scheduler service,” on page 1095

- **Asynchronous beans**

An asynchronous bean is a Java object or enterprise bean that can be executed asynchronously by a J2EE application, using the J2EE context of the bean’s creator. These beans also can run with copies of other J2EE contexts. For example:

- Internationalization context
- Application profiles
- Work areas
- Access intent policies

Asynchronous beans enable the construction of stateful, “active” J2EE applications. These applications address a segment of the application space that J2EE has not previously addressed (that is, advanced applications that require application threading, active agents within a server application, or distributed monitoring capabilities). See Chapter 18, “Using asynchronous beans,” on page 1065.

- **Object pools**

Objects are frequently pooled by Java applications in order to avoid the cost of creating new Java objects and the associated garbage collection delays that result when these objects are reclaimed after use. An object pool keeps a number of pre-allocated objects on behalf of its users. Applications can get an object from the pool, use it, and later return it to the pool. This allows the individual object instances to be reused and effectively limits the amount of garbage generated by the application. Chapter 19, “Using object pools,” on page 1085

- **Startup beans**

Startup beans are stateful session beans that enable J2EE applications to execute business logic when an application starts or stops. The startup bean is loaded when the application starts. The start() method is then invoked on the bean’s remote interface. This method can execute any business logic needed by the application at start time. Similarly, the bean’s stop() method is called on the instance when the application is stopped and can execute any business logic needed by the application at stop time. See Chapter 20, “Using startup beans,” on page 1093

Startup beans are especially useful when used in combination with asynchronous beans to develop an active J2EE server application.

- **Business Rule Beans (BRBeans)**

Business Rule Beans are used to separate business rules from an application's core behavior, allowing the application code to remain intact and untouched even as business practices change. Each business rule is represented by an entity bean that persistently stores information related to that rule. Each business rule is assigned an appropriate rule name and stored in a rule folder. The application developer identifies "points of variability" within an application and codes trigger points at these locations. These trigger points invoke one or more business rules. See Chapter 17, "Using Business Rule Beans," on page 1005

- **Transactions**

IBM WebSphere Application Server applications can use transactions to coordinate multiple updates to resources as atomic units (as indivisible units of work) such that all or none of the updates are made permanent. The way that applications use transactions depends on the type of application component, as follows:

- A session bean can either use container-managed transactions (where the bean delegates management of transactions to the container) or bean-managed transactions (where the bean manages transactions itself)
- Entity beans use container-managed transactions
- Web components (servlets) use bean-managed transactions

The product is a transaction manager that supports the coordination of resource managers through their XAResource interface and participates in distributed global transactions with other OTS 1.2 compliant transaction managers (for example, J2EE 1.3 application servers). Applications can also be configured to interact with databases, JMS queues, and JCA connectors through their local transaction support when distributed transaction coordination is not required.

Resource managers that offer transaction support can be categorized into those that support two-phase coordination (by offering an XAResource interface) and those that support only one-phase coordination (for example through a LocalTransaction interface). The IBM WebSphere Application Server transaction support provides coordination, within a transaction, for any number of two-phase capable resource managers. It also enables a single one-phase capable resource manager to be used within a transaction in the absence of any other resource managers, although a WebSphere transaction is not necessary in this case. With the Last Participant Support of Enterprise Extensions, you can coordinate the use of a single one-phase commit (1PC) capable resource with any number of two-phase commit (2PC) capable resources in the same global transaction. At transaction commit, the two-phase commit resources are prepared first using the two-phase commit protocol, and if this is successful the one-phase commit-resource is then called to `commit(one_phase)`. The two-phase commit resources are then committed or rolled back depending on the response of the one-phase commit resource.

The ActivitySession service of Enterprise Extensions provides an alternative unit-of-work (UOW) scope to that provided by global transaction contexts. It is a distributed context that can be used to coordinate multiple one-phase resource managers. The product EJB container and deployment tooling support ActivitySessions as an extension to the J2EE programming model. Enterprise beans can be deployed with lifecycles that are influenced by ActivitySession context, as an alternative to transaction context. An application can then interact with a resource manager through its LocalTransaction interface for the period of a client-scoped ActivitySession rather than just the duration of an EJB method.

- **Naming**

Naming clients use **Naming Services** primarily to access objects, such as EJB homes, associated with applications installed on IBM WebSphere Application Server. Objects are made available to clients by being bound into a name space. A name space is under the control of a name server. In this product, there are potentially many name servers, and the name spaces controlled by the various name servers are federated together to form the view of a single name space. Each name server presents the same logical view of the federated name spaces.

Name servers provided by this product are a CORBA CosNaming implementation. IBM WebSphere Application Server provides a CosNaming JNDI plug-in which enables clients to access the name servers through the JNDI interface. Clients to EJB applications typically use JNDI to perform Naming operations. Clients may access the name servers directly through the CORBA programming model. The

CosNaming interface is part of the CORBA programming model. CORBA clients which need to access EJB homes or some other objects bound to the name space would typically use the CORBA CosNaming interface to perform Naming operations.

- **Dynamic cache**

Dynamic cache improves application performance by caching outputs and contents of outputs of servlets, JavaServer Pages (JSP) files, Web services, and commands. On subsequent client requests to the same applications, dynamic cache intercepts these calls and responds by serving the output or the contents of output from the cache.

Dynamic cache in this product version includes:

- **Servlet/JSP files caching**

This caches output of dynamic servlets and JSP files by working with the Java virtual machine of the application server by intercepting calls to service methods and serving Web pages from the cache. This improves server response time, throughput and scalability.

- **Command caching**

Commands that are written to the Command Architecture encapsulate business logic tasks and provide a standard way to invoke the business logic request. Command objects need to implement CacheableCommand interface instead of TargetableCommand interface to cache. Like in servlets and JSP caching, requests to execute business logic in the command is intercepted by the cache. If a command with the same request attributes are available in cache, output properties are copied from the cached instance to the requested instance and returned without executing the business logic again.

- **Web Services caching**

Web service responses can be cached just like servlet and JSP results. These requests are intercepted and the cache ID is computed based on how the cache ID rules are specified in the cache policy. A hash of the whole SOAPEnvelope can be used as a cache ID or it can be parsed and the service name, operation name and parameter names to these operations used as cache ID. If a cache entry is not found for the computed cache ID, the request is forwarded to the SOAP engine and the result is cached.

- **Edge Side Include caching**

This provides the ability to cache, assemble and deliver dynamic web pages at the edge of the enterprise network. Edge Side Includes (ESI) is a simple markup language which enables dynamic web pages (which by themselves are not so cache efficient) to be broken down into cacheable fragments. These fragments are then cached on the edge of the network and assembled into a single page upon user requests.

- **Distributed caching**

Cache contents can be shared and replicated among servers through dynamic caching by using an underlying JMS based message broker system, DRS (Data Replication Service). Sharing characteristics of individual cache entries are configured using the cache policy specification.

- **User profiles**

Managing **user profiles** allows a company to maintain database tables containing fields for demographic data of individual customers or other users on the company system. For example, when a user repeatedly logs onto a Web site that supports user profiles, the Web site can display headlines and advertising tailored to the shopping preferences of that user. The site can address the user by his or her logon name. User profile API is deprecated in the current release.

Process choreographer

Process choreographer is a powerful enterprise workflow tool that supports running Business Process Execution Language (BPEL) processes in a J2EE environment. These processes can be used to integrate J2EE resources, Web services, and activities that require human interaction.

For more information, refer to process choreographer documentation.

Assembly tools

Assembly is an activity in which you package code components into "modules" that comply with the J2EE specification. You define configurations for the modules, in the form of XML documents known as deployment descriptors. The modules can be assembled into an enterprise application (EAR) file and deployed on a server.

5.1 + The Application Server Toolkit provides a graphical interface tool, called *Assembly Toolkit*, for assembly processes. See "Chapter 29, "Assembling applications with the Assembly Toolkit," on page 1391." For a discussion of archives and Web components supported by the Assembly Toolkit, see "Archive support in Version 5" on page 1393."

Deployment

Deployment involves placing applications onto application servers and running the applications. The main tasks include:

1. Installing application files onto an application server.
2. Configuring the application for the particular operational environment.
3. Starting the newly deployed application.

Information on these tasks is available from "Chapter 30, "Deploying and managing applications," on page 1409." The information describes how to deploy applications using the WebSphere Application Server administrative console. You can also deploy applications using the wsadmin tool, which provides deployment capabilities identical to those available using the administrative console.

Packaging and class loading

You can package your business logic as a Java 2 Platform, Enterprise Edition (J2EE) application enterprise archive (EAR) file or as an enterprise bean (EJB) or Web module for deployment to WebSphere Application Server. You must also consider the class loading relationships among modules.

Uninstalling and redeploying applications

At some point, you will need to uninstall your deployed applications. Or you might need to update your applications and deploy them again. You might be able to use hot deployment and dynamic reloading, where you do not need to restart the application server (or the application in some cases) after deploying an updated application.

Chapter 2. Using Web applications

A developer creates the files comprising a Web application, and then assembles the Web application components into a Web module. Next, the deployer (typically the developer in a unit-testing environment or the administrator in a production environment) installs the Web application on the server.

1. **(Optional)** Migrate existing Web applications to run in the new version of WebSphere.
2. Design the Web application and develop its code artifacts: Servlets, JavaServer Pages (JSP) files, and static files, as for example, images and Hyper Text Markup Language (HTML) files. See the "Resources for learning" article for links to design documentation.
3. **(Optional)** Implement JavaScript within JSP tags using the Bean Scripting Framework (BSF).

5.1+ Support in the JSP Engine for the Bean Scripting Framework is deprecated with WebSphere Application Server 5.1.

4. Develop the Web application, using WebSphere Application Server extensions to enhance its functionality.
5. Assemble the Web application into a Web module using the Assembly Toolkit. Web module assembly properties might include the ability to:
 - Configure servlet page lists.
 - Configure servlet filters.
 - Serve servlets by class name.
 - Enable file serving.
6. Deploy the Web module or application module that contains the Web application.
Following deployment, you might find it handy to use the tool that enables batch compiling of the JSP files for quicker initial response times.
7. **(Optional)** Troubleshoot your Web application. (See "Web Container troubleshooting tips" in the Information center.)
8. **(Optional)** Modify the default Web container configuration in the application server in which you deployed the Web module or application module containing the Web application.
9. **(Optional)** Manage the deployed Web application.

Web applications

A Web application is comprised of one or more related servlets, JavaServer Pages technology (JSP files), and Hyper Text Markup Language (HTML) files that you can manage as a unit.

The files in a Web application are related in that they work together to perform a business logic function.

For example, one of the WebSphere Application Server samples is a *Simple Greeting* Web application. (See "Samples Gallery" in the Information Center.) This application, comprised of a servlet and Web pages, greets new users when the application is accessed.

The Web application is a concept supported by the Java Servlet Specification. Web applications are typically packaged as .war files.

web.xml file

The web.xml file provides configuration and deployment information for the Web components that comprise a Web application. Examples of Web components are servlet parameters, servlet and JavaServer Pages (JSP) definitions, and Uniform Resource Locators (URL) mappings.

The servlet 2.3 specification dictates the format of the web.xml file, which makes this file portable among Java Two Enterprise Edition (J2EE) compliant products.

Location

The `web.xml` file must reside in the `WEB-INF` directory under the context of the hierarchy of directories that exist for a Web application. For example, if the application is `client.war`, then the `web.xml` file is placed in the `install_root/client war/WEB-INF` directory.

Usage notes

- Is this file read-only?

No

- Is this file updated by a product component?

This file is updated by the Assembly Toolkit.

- If so, what triggers its update?

The Assembly Toolkit updates the `web.xml` file when you assemble Web components into a Web module, or when you modify the properties of the Web components or the Web module.

- How and when are the contents of this file used?

WebSphere Application Server functions use information in this file during the configuration and deployment phases of Web application development.

Sample file entry

```
<?xml version="1.0" encoding="UTF-8"?>
<!DOCTYPE web-app PUBLIC "-//Sun Microsystems, Inc.
    //DTD Web Application 2.3//EN"
    "http://java.sun.com/j2ee/dtds/web-app_2_3.dtd">

<web-app id="WebApp_1">
  <display-name>Persistence Manager Web Client</display-name>
  <description>Persistence Manager Web Client</description>
  <servlet id="Servlet_1">
    <servlet-name>CustomerLocalServlet</servlet-name>
    <description>Local Customer Servlet</description>
    <servlet-class>CustomerLocalServlet</servlet-class>
  </servlet>
  <servlet id="Servlet_2">
    <servlet-name>CustomerServlet</servlet-name>
    <description>Remote Customer Servlet</description>
    <servlet-class>CustomerServlet</servlet-class>
  </servlet>
  <servlet id="Servlet_3">
    <servlet-name>CreditCardServlet</servlet-name>
    <description>Credit Card Servlet - PM Verification</description>
    <servlet-class>CreditCardServlet</servlet-class>
  </servlet>
  <servlet-mapping id="ServletMapping_1">
    <servlet-name>CustomerLocalServlet</servlet-name>
    <url-pattern>/CustomerLocal</url-pattern>
  </servlet-mapping>
  <servlet-mapping id="ServletMapping_2">
    <servlet-name>CustomerServlet</servlet-name>
    <url-pattern>/Customer</url-pattern>
  </servlet-mapping>
  <servlet-mapping id="ServletMapping_3">
    <servlet-name>CreditCardServlet</servlet-name>
    <url-pattern>/CreditCard</url-pattern>
  </servlet-mapping>
  <welcome-file-list id="WelcomeFileList_1">
    <welcome-file>index.html</welcome-file>
  </welcome-file-list>
  <security-role id="SecurityRole_1">
    <description>Everyone role</description>
    <role-name>Everyone Role</role-name>
  </security-role>
  <security-role id="SecurityRole_2">
```

```

        <description>AllAuthenticated role</description>
        <role-name>All Role</role-name>
    </security-role>
    <security-role id="SecurityRole_3">
        <description>Deny all access role</description>
        <role-name>DenyAllRole</role-name>
    </security-role>
</web-app>

```

Migrating Web application components

Supported open specification levels in WebSphere Application Server Version 5 are documented in article, "Migrating," in the Information Center.

Migration of Web applications deployed in WebSphere Application Server Version 4.x is not necessary; version 2.2 of the servlet specification and version 1.1 of the JavaServerPages (JSP) specification are still supported. However, where there are behavioral differences between the Java Two Enterprise Edition (J2EE) 1.2 and J2EE 1.3 specifications, bear in mind that J2EE 1.3 specifications are implemented in WebSphere Application Server Version 5 and will override any J2EE 1.2 behaviors.

Servlet migration might be a concern if your application:

- implements a WebSphere internal servlet to bypass a WebSphere Application Server Version 4.x single application path restriction.
- extends a PageListServlet that relies on configuration information in the servlet configuration XML file.
- uses a servlet to generate Hyper Text Markup Language (HTML) output.
- calls the `response.sendRedirect()` method for a servlet using the `encodeRedirectURL` function or executing within a non-context root.

JSP migration might be a concern if your application references JSP page implementation classes in unnamed packages, or if you install WebSphere Application Server Version 4.x EAR files (deployed in Version 4.x with the JSP Precompile option), in Version 5.

Follow these steps if migration issues apply to your Web application:

1. Use WebSphere Application Server Version 5 package names for any WebSphere Application Server Version 4.x internal servlets, which are implemented in your application.

In WebSphere Application Server Version 4.x, Web modules with a context root setting of / are not supported. Accessing Web modules with this root context results in HTTP 404 - File not Found errors.

To bypass the errors, and to enable the serving of static files from the root context, WebSphere Application Server Version 4.x users are advised to add the servlet class, `com.ibm.servlet.engine.webapp.SimpleFileServlet`, to their Web module.

The Version 4.x single path limitation does not exist in Version 5. However, users who choose to use the `com.ibm.servlet.engine.webapp.SimpleFileServlet` in Version 5 must do one of the following:

- Rename `com.ibm.servlet.engine.webapp.SimpleFileServlet` to `com.ibm.ws.webcontainer.servlet.SimpleFileServlet`.
- Open a Web deployment descriptor editor in the Assembly Toolkit and select **File serving enabled** on the **Extensions** tab.

The following list identifies the other internal servlets affected by the Version 5 package name change:

- DefaultErrorReporter
- AutoInvoker

Use the Version 5 package name, `com.ibm.ws.webcontainer.servlet.servlet class name` for these servlets.

2. Use the WASPostUpgrade migration tool to migrate servlets that extend PageListServlet and rely on configuration information in the associated XML servlet configuration file. In Version 4.x, the XML servlet configuration file provides configuration data for page lists and augments servlet configuration

information. This file is named as either *servlet_class_name*.servlet or *servlet_name*.servlet, and is stored in the same directory as the servlet class file.

The XML servlet configuration file is not supported in WebSphere Application Server Version 5.

3. Set a content type if your servlet generates Hyper Text Markup Language (HTML) output.

The default behavior of the Web container changed in WebSphere Application Server Version 5. If the servlet developer does not specify a content type in the servlet then the container is forbidden to set one automatically. Without an explicit content type setting, the content type is set to null. The Netscape browser displays HTML source as plain text with a null content type setting.

To resolve this problem, do one of the following:

- Explicitly set a content type in your servlet.
- Open a Web deployment descriptor editor in the Assembly Toolkit and select **Automatic Response Encoding enabled** on the **Extensions** tab.

4. Set the Java environment variable, `com.ibm.websphere.sendredirect.compatibility`, to **true** if you want your URLs interpreted relative to the application root.

The default value of the Java environment variable `com.ibm.websphere.sendredirect.compatibility` changed in WebSphere Application Server Version 5. In Version 4, the default setting of this variable is true. In Version 5, the setting is false.

When this variable is set to false, if a URL has a leading slash, the URL is interpreted relative to the Web module/application root. However, if the URL does not have a leading slash, it is interpreted relative to the Web container root (also known as the Web server document root). Therefore, if an application has a WAR file that has a context root of `myPledge_app` and a servlet that has a servlet mapping of `/Intranet/`, a JSP file in the WAR file cannot access the servlet when its `encodeRedirectURL` is set to `/Intranet/myPledge`. The JSP file can access the servlet if the `encodeRedirectURL` is set to `myPledge_app/Intranet/myPledge`, or if the `com.ibm.websphere.sendredirect.compatibility` variable is set to true.

See "Configuring JVM sendRedirect calls to use context root" in the Information Center.

5. Use the WASPostUpgrade migration tool to migrate WebSphere Version 4.x enterprise applications to Version 5.

Note: The WebSphere Application Server Version 4.x JSP page implementation class files are not compatible with the WebSphere Application Server Version 5 JSP container.

The WASPostUpgrade tool automatically precompiles JSP files, which ensures the JSP page implementation class files are compatible with Version 5.

If you install Version 4.x EAR files, deployed with the JSP Precompile option, in Version 5, and you choose not to follow the migration path, do one of the following:

- Select the Pre-compile JSP option in the administrative console Install New Application window.
See article [Installing a new application](#) for more information.
- Specify the `-preCompileJSPs` option when using the `wsadmin` tool.

6. Import your classes if your application uses unnamed packages.

Section 8.2 of the JSP 1.2 specification states:

The JSP container creates a JSP page implementation class for each JSP page. The name of the JSP page implementation class is implementation dependent. The JSP page implementation object belongs to an implementation-dependent named package. The package used may vary between one JSP and another, so minimal assumptions should be made. The unnamed package should not be used without an explicit *import* of the class.

For example, if `myBeanClass` is in the unnamed package, and you reference it in a `jsp:useBean` tag, then you must explicitly import `myBeanClass` with the page directive `import` attribute, as shown in the following example:

```
<%@page import="myBeanClass" %>
.
.
.
<jsp:useBean id="myBean" class="myBeanClass" scope="session"/>
```

In WebSphere Application Server Version 5, the JSP engine creates JSP page implementation classes in the `org.apache.jsp` package. If a class in the unnamed package is not explicitly imported, then the `javac` compiler assumes the class is in package `org.apache.jsp`, and the compilation fails.

The *Incompatibilities* section of the version 1.4.Java 2 Platform, Standard Edition (J2SE) documentation states:

The compiler now rejects import statements that import a type from the unnamed namespace. Previous versions of the compiler would accept such import declarations, even though they were arguably not allowed by the language (because the type name appearing in the import clause is not in scope). The specification is being clarified to state clearly that you cannot have a simple name in an import statement, nor can you import from the unnamed namespace.

To summarize, the syntax:

```
import SimpleName;
```

is no longer legal. Nor is the syntax

```
import ClassInUnnamedNamespace.Nested;
```

which would import a nested class from the unnamed namespace.

To fix such problems in your code, move all of the classes from the unnamed namespace into a named namespace.

See "Resources for learning" for links to the J2SE, JSP, and Servlet specification documentation.

Default Application

The IBM WebSphere Application Server provides a default configuration that allows administrators to easily verify that the Application Server is running. When the product is installed, it includes an application server called *server1* and an enterprise application called *Default Application*.

Default Application contains a Web Module called *DefaultWebApplication* and an enterprise bean JAR file called *Increment*. The *Default Application* provides a number of servlets, described below. These servlets are available in the product.

For additional code examples, visit the Samples Gallery. Learn how to locate and install the Samples Gallery by viewing the Samples Gallery reference page in the Information Center.

The URL for accessing Samples is: `http://localhost:9080/WSamples/`

Snoop

Use the Snoop servlet to retrieve information about a servlet request. This servlet returns the following information:

- Servlet initialization parameters
- Servlet context initialization parameters
- URL invocation request parameters
- Preferred client locale
- Context path
- User principal
- Request headers and their values
- Request parameter names and their values
- HTTPS protocol information
- Servlet request attributes and their values
- HTTP session information
- Session attributes and their values

The Snoop servlet includes security configuration so that when WebSphere Security is enabled, clients must supply a user ID and password to execute the servlet.

The URL for the Snoop servlet is: `http://localhost:9080/snoop/`.

HelloHTML

Use the HelloHTML pervasive servlet to exercise the PageList support provided by the WebSphere Web container. This servlet extends the PageListServlet, which provides APIs that allow servlets to call other Web resources by name or, when using the *Client Type detection* support, by type.

You can invoke the Hello servlet from an HTML browser, speech client, or most Wireless Application Protocol (WAP) enabled browsers using the URL: `http://localhost:9080/HelloHTML.jsp`.

HitCount

Use the HitCount Demonstration application to demonstrate incrementing a counter using a variety of methods, including:

- A servlet instance variable
- An HTTP session
- An enterprise bean

You can instruct the servlet to execute any of these methods within a transaction that you can omit or roll back. If the transaction is committed, the counter is incremented. If the transaction is rolled back, the counter is not incremented.

The enterprise bean method uses a Container- Managed Persistence enterprise bean that persists the counter value to a Cloudscape database. This enterprise bean is configured to use the Default Datasource, which is set to the DefaultDB database.

When using the enterprise bean method, you can instruct the servlet to look up the enterprise bean, either in the WebSphere global namespace, or in the namespace local to the application.

The URL for the HitCount application is: `http://localhost:9080/HitCount.jsp`.

Servlets

Servlets are Java programs that use the Java Servlet Application Programming Interface (API). You must package servlets in a Web ARchive (WAR) file or Web module for deployment to the application server. *Servlets* run on a Java-enabled Web server and extend the capabilities of a Web server, similar to the way applets run on a browser and extend the capabilities of a browser.

Servlets can support dynamic Web page content, provide database access, serve multiple clients at one time, and filter data.

For the purposes of IBM WebSphere Application Server, discussions of servlets focus on Hyper Text Transfer Protocol (HTTP) servlets, which serve Web-based clients.

Developing servlets with WebSphere Application Server extensions

Several WebSphere Application Server extensions are provided for enhancing your servlets. This task provides a summary of the extensions that you can utilize.

1. Review the supported specifications.

Create Java components, referring to the Servlet specifications from Sun Microsystems.

See Resources for learning for links to coding specifications and examples.

The application server includes its own packages that extend and add to the Java Servlet Application Programming Interface (API). These extensions and additions make it easier to manage session states, create personalized Web pages, generate better servlet error reports, and access databases. Locate the Javadoc for the application server APIs in the product `install_root\web\apidocs` directory.

All the public WebSphere Application Server APIs are located in the `com.ibm.websphere...` packages.

2. Use your favorite integrated development environment (IDE), or a text editor, to develop or migrate code artifacts that meet the specifications.
3. Test the code artifacts.

Assemble your code artifacts into a Web module using the Assembly Toolkit as a prerequisite to deploying the code to the application server.

Application lifecycle listeners and events

Application lifecycle listeners and events, now part of the Servlet API, enable you to notify interested listeners when servlet contexts and sessions change. For example, you can notify users when attributes change and if sessions or servlet contexts are created or destroyed.

The lifecycle listeners give the application developer greater control over interactions with `ServletContext` and `HttpSession` objects. Servlet context listeners manage resources at an application level. Session listeners manage resources associated with a series of requests from a single client. Listeners are available for lifecycle events and for attribute modification events. The listener developer creates a class that implements the `javax.listener` interface, corresponding to the desired listener functionality.

At application startup time, the container uses introspection to create an instance of your listener class and registers it with the appropriate event generator.

When a servlet context is created, the `contextInitialized` method of your listener class is invoked, which creates the database connection for the servlets in your application to use, if this context is for your application.

When the servlet context is destroyed, your `contextDestroyed` method is invoked, which releases the database connection, if this context is for your application.

Listener classes for servlet context and session changes

The following methods are defined as part of the `javax.servlet.ServletContextListener` interface:

- `void contextInitialized(ServletContextEvent)` - Notification that the Web application is ready to process requests.

Place code in this method to see if the created context is for your Web application and if it is, allocate a database connection and store the connection in the servlet context.

- `void contextDestroyed(ServletContextEvent)` - Notification that the servlet context is about to shut down.

Place code in this method to see if the created context is for your Web application and if it is, close the database connection stored in the servlet context.

Two new listener interfaces are defined as part of the `javax.servlet` package:

- `ServletContextListener`
- `ServletContextAttributeListener`

One new filter interface is defined as part of the `javax.servlet` package:

- `FilterChain` interface - methods: `doFilter()`

Two new event classes are defined as part of the `javax.servlet` package:

- ServletContextEvent
- ServletContextAttributeEvent

Three new listener interfaces are defined as part of the javax.servlet.http package:

- HttpSessionListener
- HttpSessionAttributeListener
- HttpSessionActivationListener

One new event class is defined as part of the javax.servlet.http package:

- HttpSessionEvent

Example: com.ibm.websphere.DBConnectionListener.java

The following example shows how to create a servlet context listener:

```
package com.ibm.websphere;

import java.io.*;
import javax.servlet.*;

public class DBConnectionListener implements ServletContextListener
{
    // implement the required context init method
    void contextInitialized(ServletContextEvent sce)
    {
    }

    // implement the required context destroy method
    void contextDestroyed(ServletContextEvent sce)
    {
    }
}
```

Servlet filtering

Servlet filtering is an integral part of the Servlet 2.3 API. Servlet filtering provides a new type of object called a *filter* that can transform a request or modify a response.

You can chain filters together so that a group of filters can act on the input and output of a specified resource or group of resources.

Filters typically include logging filters, image conversion filters, encryption filters, and Multipurpose Internet Mail Extensions (MIME) type filters (functionally equivalent to the servlet chaining). Although filters are not servlets, their lifecycle is very similar.

Filters are handled in the following manner:

- The Web container determines whether it needs to construct a FilterChain containing the LoggingFilter for the requested resource.
The FilterChain begins with the invocation of the LoggingFilter and ends with the invocation of the requested resource.
- If other filters need to go in the chain, the Web container places them after the LoggingFilter and before the requested resource.
- The Web container then instantiates and initializes the LoggingFilter (if it was not done previously) and invokes its doFilter(FilterConfig) method to start the chain.
- The LoggingFilter preprocesses the request and response objects and then invokes the filter chain doFilter(ServletRequest, ServletResponse) method.

This method passes the processing to the next resource in the chain (in this case, the requested resource).

- Upon return from the filter chain `doFilter(ServletRequest, ServletResponse)` method, the `LoggingFilter` performs post-processing on the request and response object before sending the response back to the client.

Filter, FilterChain, FilterConfig classes for servlet filtering

The following interfaces are defined as part of the `javax.servlet` package:

- `Filter` interface - methods: `doFilter()`, `getFilterConfig()`, `setFilterConfig()`
- `FilterChain` interface - methods: `doFilter()`
- `FilterConfig` interface - methods: `getFilterName()`, `getInitParameter()`, `getInitParameterNames()`, `getServletContext()`

The following classes are defined as part of the `javax.servlet.http` package:

- `HttpServletRequestWrapper` - methods: See the Servlet 2.3 Specification
- `HttpServletResponseWrapper` - methods: See the Servlet 2.3 Specification

Example: `com.ibm.websphere.LoggingFilter.java`

The following example shows how to implement a filter:

```
package com.ibm.websphere;

import java.io.*;
import javax.servlet.*;

public class LoggingFilter implements Filter
{
    File _loggingFile = null;

    // implement the required init method
    public void init(FilterConfig fc)
    {
        // create the logging file
        xxx;
    }

    // implement the required doFilter method...this is where most of
    the work is done
    public void doFilter(ServletRequest request,
        ServletResponse response, FilterChain chain)
    {
        try
        {
            // add request info to the log file
            synchronized(_loggingFile)
            {
                xxx;
            }

            // pass the request on to the next resource in the chain
            chain.doFilter(request, response);
        }
        catch (Throwable t)
        {
            // handle problem...
        }
    }

    // implement the required destroy method
    public void destroy()
    {
        // make sure logging file is closed
        _loggingFile.close();
    }
}
```

Configuring page list servlet client configurations

You can define PageListServlet configuration information in the IBM Web Extensions file. The IBM Web Extensions file is created and stored in the Web Applications archive (WAR) file by the Assembly Toolkit.

To configure and implement page lists:

1. To configure page list information, use the Add Markup Language entry dialog of the Assembly Toolkit. On the **Servlets** tab of a Web deployment descriptor editor, select a servlet and click **Add** under **WebSphere Extensions**.
2. Add the `callPage()` method to your servlet to invoke a JavaServer Page (JSP) file in response to a client request.

The PageListServlet has a `callPage()` method that invokes a JSP file in response to the HTTP request for a page in a page list. The `callPage()` method can be invoked in one of the following ways:

- `callPage(String pageName, HttpServletRequest request, HttpServletResponse response)`

where the method arguments are:

pageName

A page name defined in the PageListServlet configuration

request

The HttpServletRequest object

response

The HttpServletResponse object

- `callPage(String mName, String pageName, HttpServletRequest request, HttpServletResponse response)`

where the method arguments are:

mName A markup language type

pageName

A page name defined in the PageListServlet configuration

request

The HttpServletRequest object

response

The HttpServletResponse object

3. Use the PageList Servlet client type detection support to determine the markup language type a calling client requires for the response.

Page lists

Page lists allow you to avoid hard-coding URLs in servlets and JSP files. A page list specifies the location where a request is to be forwarded, but automatically customizes that location depending on the MIME type of the servlet. Use these properties to specify a markup language and an associated MIME type. For the given MIME type, you also specify a set of pages to invoke.

WebSphere Application Server supplies the PageListServlet servlet, which you can use to call a JavaServer Pages (JSP) file by name based on the configuration data in the `client_types.xml` file. This file maps a JSP file to a Uniform Resource Identifier (URI). When the URI is invoked, it specifies another JSP file in a Web module. This support allows you to access multiple Uniform Resource Locators (URLs) without hard-coding them in your servlets.

You can also logically group page lists according to the markup language type, such as, Hypertext Markup Language (HTML) or Wireless Markup Language (WML). This allows applications that use servlets to extend the PageListServlet servlet, to call JSP files which return the proper markup-language type for the client request. For example, a request that originates from a PDA device requires WML data. The application server sends the request to a servlet that extends the PageListServlet servlet, and the servlet calls a JSP file that returns a WML response.

Client type detection support

In addition to providing the page list mapping capability, the `PageListServlet` also provides *Client Type Detection* support. A servlet determines the markup language type that a calling client needs in the response, using the configuration information in the `client_types.xml` file.

Client type detection support allows a servlet, extending the `PageListServlet`, to call an appropriate JavaServer Pages (JSP) file. The servlet invokes the `callPage()` method, which calls a JSP file based on the markup-language type of the request.

client_types.xml

The `client_types.xml` file provides client type detection support for servlets extending `PageListServlet`. Using the configuration data in the `client_types.xml` file, servlets can determine the language type that calling clients require for the response.

The client type detection support allows servlets to call appropriate JavaServer Pages (JSP) files with the `callPage()` method. Servlets select JSP files based on the markup-language type of the request.

Servlets must use the following version of the `callPage()` method to determine the markup language type required by the client:

```
callPage(String mlName, String pageName, HttpServletRequest request,
         HttpServletResponse response)
```

where the arguments are:

- `mlName` - a markup language type
- `pageName` - a page name defined in the `PageListServlet` configuration
- `request` - the `HttpServletRequest` object
- `response` - the `HttpServletResponse` object

Review the Extending `PageListServlet` code example to see how the `callPage()` method is invoked by a servlet.

In the example, the client type detection method, `getMLTypeFromRequest(HttpServletRequest request)`, provided by the `PageListServlet`, inspects the `HttpServletRequest` object request headers, and searches for a match in the `client_types.xml` file.

The client type detection method does the following:

- Uses the input `HttpServletRequest` and the `client_types.xml` file, to check for a matching HTTP request name and value.
- Returns the markup-language value configured for the `<client-type>` element, if a match is found. If multiple matches are found, this method returns the markup-language for the first `<client-type>` element for which a match is found.
- If no match is found, returns the value of the markup-language for the default page defined in the `PageListServlet` configuration.

Location

The `client_types.xml` file is located in the `install_root/properties` directory.

Usage notes

- Is this file read-only?
No
- Is this file updated by a product component?
No
- If so, what triggers its update?
This file is created and updated manually by users.

- How and when are the contents of this file used?

Servlets, extending `PageListServlet`, use this file to determine the language type that calling clients require for the response.

Sample file entry

```
<?xml version="1.0" >
<!DOCTYPE clients [
<!ELEMENT client-type (description, markup-language,request-header+)>
<!ELEMENT description (#PCDATA)>
<!ELEMENT markup-language (#PCDATA)>
<!ELEMENT request-header (name, value)>
<!ELEMENT clients (client-type+)>
<!ELEMENT name (#PCDATA)>
<!ELEMENT value (#PCDATA)>]>
<clients>
  <client-type>
    <description>IBM Speech Client</description>
    <markup-language>VXML</markup-language>
    <request-header>
      <name>user-agent</name>
      <value>IBM VoiceXML pre-release version 000303</value>
    </request-header>
    <request-header>
      <name>accept</name>
      <value>text/vxml</value>
    </request-header>
  </client-type>
  <client-type>
    <description>WML Browser</description>
    <markup-language>WML</markup-language>
    <request-header>
      <name>accept</name>
      <value>text/x-wap.wml</value>
    </request-header>
    <request-header>
      <name>accept</name>
      <value>text/vnd.wap.xml</value>
    </request-header>
  </client-type>
</clients>
```

Example: Extending PageListServlet

The following example shows how a servlet extends the `PageListServlet` class and determines the markup-language type required by the client. The servlet then uses the `callPage()` method to call an appropriate JavaServer Pages (JSP) file. In this example, the JSP file that provides the the correct markup-language for the response is *Hello.page*.

```
public class HelloPervasiveServlet extends PageListServlet implements Serializable
{
  /*
  * doGet -- Process incoming HTTP GET requests
  */
  public void doGet(HttpServletRequest request, HttpServletResponse response)
  throws IOException, ServletException
  {
    // This is the name of the page to be called:
    String pageName = "Hello.page";

    // First check if the servlet was invoked with a queryString that contains
    // a markup-language value.
    // For example, if this is how the servlet is invoked:
    // http://localhost/servlets/HelloPervasive?mlname=VXML
    // then use the following method:
    String mlname= getMLNameFromRequest(request);
```

```

// If no markup language type is provided in the queryString,
// then try to determine the client
// Type from the request, and use the markup-language name configured in
// the client_types.xml file.
if (mlName == null)
{
    mlName = getMLTypeFromRequest(request);
}
try
{
    // Serve the request page.
    callPage(mlName, pageName, request, response);
}
catch (Exception e)
{
    handleError(mlName, request, response, e);
}
}
}

```

autoRequestEncoding and autoResponseEncoding

Two new WebSphere Application Server extensions are available in Version 5, `autoRequestEncoding` and `autoResponseEncoding`.

In WebSphere Application Server Version 5, the Web container no longer automatically sets request and response encodings, and response content types. Programmers are expected to set these values using available methods in the Servlet 2.3 Specification. If programmers choose not to use the character encoding methods, they can specify the `autoRequestEncoding` and `autoResponseEncoding` extensions, which enable the application server to set the encoding values and content type.

The values of the `autoRequestEncoding` and `autoResponseEncoding` extensions are either `true` or `false`. The default value for both extensions is `false`. If the value is `false` for both `autoRequestEncoding` and `autoResponseEncoding`, then the request and response character encoding is set to the Servlet 2.3 Specification default, which is ISO-8859-1. Also, if the value is set to `false` for a response, the Web container cannot set a response content type.

Use the Assembly Toolkit to change the default values for the `autoRequestEncoding` and `autoResponseEncoding` extensions.

Review the `autoRequestEncoding` and `autoResponseEncoding` encoding examples for a description of Web container behavior when these values are set to `true`.

Examples: autoRequestEncoding and autoResponseEncoding encoding examples

The default value of the `autoRequestEncoding` and `autoResponseEncoding` extensions is `false`, which means that both the request and response character encoding is set to the Servlet 2.3 Specification default of ISO-8859-1. Different character encodings are possible if the client defines character encoding in the request header, or if the code includes the `setCharacterEncoding(String encoding)` method. Also, if the value is set to `false` for a response, the Web container cannot set a response content type.

If the `autoRequestEncoding` value is set to `true`, and the client did not specify character encoding in the request header, and the code does not include the `setCharacterEncoding(String encoding)` method, the Web container tries to determine the correct character encoding for the request parameters and data.

The Web container performs each step in the following list until a match is found:

- Looks at the character set (charset) in the *Content-Type* header.
- Attempts to map the servers locale to a character set using defined properties.
- Attempts to use the `DEFAULT_CLIENT_ENCODING` system property, if one is set.

- Uses the ISO-8859-1 character encoding as the default.

If the `autoResponseEncoding` value is set to `true`, and the client did not specify character encoding in the request header, and the code does not include the `setCharacterEncoding(String encoding)` method, the Web container does the following:

- Attempts to determine the response content type and character encoding from information in the request header.
- Uses the ISO-8859-1 character encoding as the default.

JavaServer Pages files

JavaServer Pages (JSP) files are application components coded to the Sun Microsystems JavaServer Pages (JSP) Specification. JSP files enable the separation of the Hypertext Markup Language (HTML) code from the business logic in Web pages so that HTML programmers and Java programmers can more easily collaborate in creating and maintaining pages.

The IBM extensions to the JSP Specification include JSP tags that resemble HTML tags. These JSP tags make it easy for HTML authors to add the power of Java technology to Web pages, without being experts in Java programming.

JSP files support a division of roles:

HTML authors

Develop JSP files that access databases and reusable Java components, such as servlets and beans.

Java programmers

Create the reusable Java components and provide the HTML authors with the component names and attributes.

Database administrators

Provide the HTML authors with the name of the database access and table information.

Developing JavaServer Pages files with WebSphere extensions

Several IBM WebSphere extensions are provided for enhancing your JavaServer Pages (JSP) files. This task provides a summary of the extensions that you can utilize.

1. Review the supported specifications.

Create Java components, referring to the JSP specifications from Sun Microsystems.

See [Resources for learning](#) for links to coding specifications and examples.

WebSphere Application Server Version 3.5 added IBM extensions to the base Application Programming Interfaces (APIs). Since the JavaServer Pages (JSP) 1.1 and JSP 1.2 Specifications are backward compatible to the JSP 1.0 Specifications, you can invoke the APIs with the IBM extensions without modification.

The extensions belong to these categories:

Syntax for variable data

Put variable fields in JSP files and have servlets and beans dynamically replace the variables with values from a database when the JSP output is returned to the browser.

Syntax for database access

Add a database connection to a Web page and then use that connection to query or update the database. You can provide the user ID and password for the database connection at request time, or you can hard code the user ID and password within the JSP file.

2. Use your favorite integrated development environment (IDE), or a text editor, to develop or migrate code artifacts that meet the specifications.
3. Test the code artifacts.
4. **(Optional)** Batch compile your JSP files if necessary.

Tag libraries

Java ServerPages (JSP) tag libraries contain classes for common tasks such as processing forms and accessing databases from JSP files.

Tag libraries encapsulate, as simple tags, core functionality common to many Web applications. The Java Standard Tag Library (JSTL) supports common programming tasks such as iteration and conditional processing, and provides tags for:

- manipulating XML documents
- supporting internationalization
- using Structured Query Language (SQL)

Tag libraries also introduce the concept of an expression language to simplify page development, and include a version of the JSP expression language.

A tag library has two parts - a Tag Library Descriptor (TLD) file and a JAR file.

tsx:dbconnect tag JavaServer Pages syntax

Use the <tsx:dbconnect> tag to specify information needed to make a connection to a Java Database Connectivity (JDBC) or an Open Database Connectivity (ODBC) database.

The <tsx:dbconnect> syntax does not establish the connection. Use the <tsx:dbquery> and <tsx:dbmodify> syntax instead to reference a <tsx:dbconnect> tag in the same JavaServer Pages (JSP) file to establish the connection.

When the JSP file compiles into a servlet, the Java processor adds the Java coding for the <tsx:dbconnect> syntax to the servlet service() method, which means a new database connection is created for each request for the JSP file.

This section describes the syntax of the <tsx:dbconnect> tag.

```
<tsx:dbconnect id="connection_id"
  userid="db_user" passwd="user_password"
  url="jdbc:subprotocol:database"
  driver="database_driver_name"
  jndiname="JNDI_context/logical_name">
</tsx:dbconnect>
```

where:

- **id**

Represents a required identifier. The scope is the JSP file. This identifier is referenced by the connection attribute of a <tsx:dbquery> tag.

- **userid**

Represents an optional attribute that specifies a valid user ID for the database that you want to access. Specify this attribute to add the attribute and its value to the request object.

Although the userid attribute is optional, you must provide the user ID. See <tsx:userid> and <tsx:passwd> for an alternative to hard coding this information in the JSP file.

- **passwd**

Represents an optional attribute that specifies the user password for the userid attribute. (This attribute is not optional if the userid attribute is specified.) If you specify this attribute, the attribute and its value are added to the request object.

Although the passwd attribute is optional, you must provide the password. See <tsx:userid> and <tsx:passwd> for an alternative to hard coding this attribute in the JSP file.

- **url** and **driver**

Represents a required attribute if you want to establish a database connection. You must provide the URL and driver.

The application server supports connection to JDBC databases and ODBC databases.

- For a JDBC database, the URL consists of the following colon-separated elements: jdbc, the subprotocol name, and the name of the database to access. An example for a connection to the Sample database included with IBM DB2 is:

```
url="jdbc:db2:sample"  
driver="COM.ibm.db2.jdbc.app.DB2Driver"
```

- For an ODBC database, use the Sun JDBC-to-ODBC bridge driver included in their Java2 Software Developers Kit (SDK) or another vendor's ODBC driver.

The url attribute specifies the location of the database. The driver attribute specifies the name of the driver to use in establishing the database connection.

If the database is an ODBC database, you can use an ODBC driver or the Sun JDBC-to-ODBC bridge. If you want to use an ODBC driver, refer to the driver documentation for instructions on specifying the database location with the url attribute and the driver name.

If you use the bridge, the url syntax is jdbc:odbc:database. An example follows:

```
url="jdbc:odbc:autos"  
driver="sun.jdbc.odbc.JdbcOdbcDriver"
```

Note: To enable the application server to access the ODBC database, use the ODBC Data Source Administrator to add the ODBC data source to the System DSN configuration. To access the ODBC Administrator, click the ODBC icon on the Windows NT Control Panel.

- **jdbcname**

Represents an optional attribute that identifies a valid context in the application server Java Naming and Directory Interface (JNDI) naming context and the logical name of the data source in that context. The Web administrator configures the context using an administrative client such as the WebSphere Administrative Console.

If you specify the jdbcname attribute, the JSP processor ignores the driver and url attributes on the <tsx:dbconnect> tag.

An empty element (such as <url></url>) is valid.

dbquery tag JavaServer Pages syntax

Use the <tsx:dbquery> tag to establish a connection to a database, submit database queries, and return the results set.

The <tsx:dbquery> tag does the following:

1. References a <tsx:dbconnect> tag in the same JavaServer Pages (JSP) file and uses the information the tag provides to determine the database URL and driver. You can also obtain the user ID and password from the <tsx:dbconnect> tag if those values are provided in the <tsx:dbconnect> tag.
2. Establishes a new connection
3. Retrieves and caches data in the results object.
4. Closes the connection and releases the connection resource.

This section describes the syntax of the <tsx:dbquery> tag.

```
<%-- SELECT commands and (optional) JSP syntax can be placed within the tsx:dbquery. --%>  
<%-- Any other syntax, including HTML comments, are not valid. --%>  
<tsx:dbquery id="query_id" connection="connection_id" limit="value" >  
</tsx:dbquery>
```

where:

- **id**

Represents the identifier of this query. The scope is the JSP file. Use id to reference the query. For example, from the <tsx:getProperty> tag, use id to display the query results.

The id is a tsx reference to the bean and can be used to retrieve the bean from the page context. For example, if id is named mySingleDBBean, instead of using:

- if (mySingleDBBean.getValue("UISEAM",0).startsWith("N"))

use:

- `com.ibm.ws.webcontainer.jsp.tsx.db.QueryResults bean = (com.ibm.ws.webcontainer.jsp.tsx.db.QueryResults)pageContext. findAttribute("mySingleDBBean"); if (bean.getValue("UISEAM",0).startsWith("N")). . .`

The bean properties are dynamic and the property names are the names of the columns in the results set. If you want different column names, use the SQL keyword for specifying an alias on the SELECT command. In the following example, the database table contains columns named FNAME and LNAME, but the SELECT statement uses the AS keyword to map those column names to FirstName and LastName in the results set:

```
Select FNAME, LNAME AS FirstName, LastName from Employee where FNAME='Jim'
```

- **connection**

Represents the identifier of a `<tsx:dbconnect>` tag in this JSP file. The `<tsx:dbconnect>` tag provides the database URL, driver name, and optionally, the user ID and password for the connection.

- **limit**

Represents an optional attribute that constrains the maximum number of records returned by a query. If this attribute is not specified, no limit is used. In such a case, the effective limit is determined by the number of records and the system caching capability.

- **SELECT command and JSP syntax**

Represents the only valid SQL command, SELECT. The `<tsx:dbquery>` tag must return a results set. Refer to your database documentation for information about the SELECT command. See other articles in this section for a description of JSP syntax for variable data and inline Java code.

dbmodify tag JavaServer Pages syntax

The `<tsx:dbmodify>` tag establishes a connection to a database and then adds records to a database table.

The `<tsx:dbmodify>` tag does the following:

1. References a `<tsx:dbconnect>` tag in the same JavaServer Pages (JSP) file and uses the information provided by that tag to determine the database URL and driver.
Note: You can also obtain the user ID and password from the `<tsx:dbconnect>` tag if those values are provided in the `<tsx:dbconnect>` tag.
2. Establishes a new connection.
3. Updates a table in the database.
4. Closes the connection and releases the connection resource.

This section describes the syntax of the `<tsx:dbmodify>` tag.

```
<%-- Any valid database update commands can be placed within the DBMODIFY tag. -->  
<%-- Any other syntax, including HTML comments, are not valid. -->  
<tsx:dbmodify connection="connection_id">  
</tsx:dbmodify>
```

where:

- **connection**

Represents the identifier of a `<tsx:dbconnect>` tag in this JSP file. The `<tsx:dbconnect>` tag provides the database URL, driver name, and (optionally) the user ID and password for the connection.

- **Database commands**

Represents valid database commands. Refer to your database documentation for details

tsx:getProperty tag JavaServer Pages syntax and examples

The `<tsx:getProperty>` tag gets the value of a bean to display in a JavaServer Pages (JSP) file.

This IBM extension of the Sun JSP `<jsp:getProperty>` tag implements all of the `<jsp:getProperty>` function and adds the ability to introspect a database bean created using the IBM extension `<tsx:dbquery>` or `<tsx:dbmodify>`.

Note: You cannot assign the value from this tag to a variable. The value, generated as output from this tag, displays in the browser window.

This section describes the syntax of the `<tsx:getProperty>` tag:

```
<tsx:getProperty name="bean_name"
  property="property_name" />
```

where:

- **name**
Represents the name of the bean declared by the `id` attribute of a `<tsx:dbquery>` syntax within the JSP file. See `<tsx:dbquery>` for an explanation. The value of this attribute is case-sensitive.
- **property**
Represents the property of the bean to access for substitution. The value of the attribute is case-sensitive and is the locale-independent name of the property.

Tag example:

```
<tsx:getProperty name="userProfile" property="username" />
<tsx:getProperty name="request" property=request.getParameter("corporation") />
```

In most cases, the value of the property attribute is just the property name. However, to access the request bean or to access a property of a property (sub property), specify the full form of the property attribute. The full form also gives you the option to specify an index for indexed properties. You can specify the optional index as a constant (such as 2), or an index like the one described in the `<tsx:repeat>` tag. Some examples using the full form of the property attribute follow:

```
<tsx:getProperty name="staffQuery" property=address(currentAddressIndex) />
<tsx:getProperty name="shoppingCart" property=items(4).price />
<tsx:getProperty name="fooBean" property=foo(2).bat(3).boo.far />
```

tsx:userid and tsx:passwd tag JavaServer Pages syntax

With the `<tsx:userid>` and `<tsx:passwd>` tags, you do not have to hard code a user ID and password in the `<tsx:dbconnect>` tag.

Use the `<tsx:userid>` and `<tsx:passwd>` tags to accept user input for the values and then add that data to the request object. You can access the request object with a JavaServer Pages (JSP) file, such as the *JSPEmployee.jsp* example that requests the database connection.

You must use `<tsx:userid>` and `<tsx:passwd>` tags within a `<tsx:dbconnect>` tag.

This section describes the syntax of the `<tsx:userid>` and `<tsx:passwd>` tags.

```
<tsx:dbconnect id="connection_id"
  <font color="red"><userid></font>
  <tsx:getProperty name="request" property=request.getParameter("userid") />
  <font color="red"></userid></font>
  <font color="red"><passwd></font>
  <tsx:getProperty name="request" property=request.getParameter("passwd") />
  <font color="red"></passwd></font>
  url="protocol:database_name:database_table"
  driver="JDBC_driver_name">
</tsx:dbconnect>
```

where:

- **<tsx:getProperty>**

Represents the syntax as a mechanism for embedding variable data.

- **userid**

Represents a reference to the request parameter that contains the user ID. You must add the parameter to the request object that passes to this JSP file. You can set the attribute and its value in the request object, using an HTML form or a URL query string to pass the user-specified request parameters.

- **passwd**

Represents a reference to the request parameter that contains the password. Add the parameter to the request object that passes to this JSP file. You can set the attribute and its value in the request object, using an HTML form or a URL query string, to pass user-specified values.

tsx:repeat tag JavaServer Pages syntax

The <tsx:getProperty> tag repeats a block of HTML tagging.

Use the <tsx:repeat> syntax to iterate over a database query results set. The <tsx:repeat> syntax iterates from the start value to the end value until one of the following conditions is met:

- The end value is reached.
- An exception is thrown.

The output of a <tsx:repeat> block is buffered until the block completes. If an exception is thrown before a block completes, no output is written for that block.

This section describes the syntax of the <tsx:repeat> tag:

```
<tsx:repeat index=name start="starting_index" end="ending_index">
</tsx:repeat>
```

where:

- **index**

Represents an optional name used to identify the index of this repeat block. The value is case-sensitive and its scope is the JSP file.

- **start**

Represents an optional starting index value for this repeat block. The default is 0.

- **end**

Represents an optional ending index value for this repeat block. The maximum value is 2,147,483,647.

If the value of the end attribute is less than the value of the start attribute, the end attribute is ignored.

Example: Combining tsx:repeat and tsx:getProperty JavaServer Pages tags

The following code snippet shows you how to code these tags:

```
<tsx:repeat>
<tr>
  <td><tsx:getProperty name="empqs" property="EMPNO" />
  <tsx:getProperty name="empqs" property="FIRSTNAME" />
  <tsx:getProperty name="empqs" property="WORKDEPT" />
  <tsx:getProperty name="empqs" property="EDLEVEL" />
</td>
</tr>
</tsx:repeat>
```

Example: tsx:dbmodify tag syntax

In the following example, a new employee record is added to a database. The values of the fields are based on user input from this JavaServer Pages (JSP) file and referenced in the database commands using the <tsx:getProperty> tag.

```

<tsx:dbmodify connection="conn" >
insert into EMPLOYEE
  (EMPNO,FIRSTNME,MIDINIT,LASTNAME,WORKDEPT,EDLEVEL)
values
('<tsx:getProperty name="request" property=request.getParameter("EMPNO") />',
'<tsx:getProperty name="request" property=request.getParameter("FIRSTNME") />',
'<tsx:getProperty name="request" property=request.getParameter("MIDINIT") />',
'<tsx:getProperty name="request" property=request.getParameter("LASTNAME") />',
'<tsx:getProperty name="request" property=request.getParameter("WORKDEPT") />',
'<tsx:getProperty name="request" property=request.getParameter("EDLEVEL") />')
</tsx:dbmodify>

```

Example: Using `tsx:repeat` JavaServer Pages tag to iterate over a results set

The `<tsx:repeat>` tag iterates over a results set. The results set is contained within a bean. The bean can be a static bean, for example, a bean created by using the IBM WebSphere Studio database wizard, or a dynamically generated bean, for example, a bean generated by the `<tsx:dbquery>` syntax. The following table is a graphic representation of the contents of a bean called, *myBean*:

	col1	col2	col3
row0	friends	Romans	countrymen
row1	bacon	lettuce	tomato
row2	May	June	July

Some observations about the bean:

- The column names in the database table become the property names of the bean. The `<tsx:dbquery>` section describes a technique for mapping the column names to different property names.
- The bean properties are indexed. For example, `myBean.get(Col1(row2))` returns May.
- The query results are in the rows. The `<tsx:repeat>` tag iterates over the rows, beginning at the start row.

The following table compares using the `<tsx:repeat>` tag to iterate over a static bean, versus a dynamically generated bean:

Static Bean Example	<tsx:repeat> Bean Example
<p>myBean.class</p> <pre>// Code to get a connection // Code to get the data Select * from myTable; // Code to close the connection</pre> <p>JSP file</p> <pre><tsx:repeat index=abc> <tsx:getProperty name="myBean" property="coll(abc)" /> </tsx:repeat></pre> <p>Notes:</p> <ul style="list-style-type: none"> • The bean (myBean.class) is a static bean. • The method to access the bean properties is myBean.get(<i>property(index)</i>). • You can omit the property index, in which case the index of the enclosing <tsx:repeat> tag is used. You can also omit the index on the <tsx:repeat> tag. • The <tsx:repeat> tag iterates over the bean properties row by row, beginning with the start row. 	<p>JSP file</p> <pre><tsx:dbconnect id="conn" userid="alice"passwd="test" url="jdbc:db2:sample" driver="COM.ibm.db2.jdbc.app.DB2Driver"> </tsx:dbconnect > <tsx:dbquery id="dynamic" connection="conn" > Select * from myTable; </tsx:dbquery> <tsx:repeat index=abc> <tsx:getProperty name="dynamic" property="coll(abc)" /> </tsx:repeat></pre> <p>Notes:</p> <ul style="list-style-type: none"> • The bean (dynamic) is generated by the <tsx:dbquery> tag and does not exist until the syntax executes. • The method to access the bean properties is dynamic.getValue(<i>"property", index</i>). • You can omit the property index, in which case the index of the enclosing <tsx:repeat> tag is used. You can also omit the index on the <tsx:repeat> tag. • The <tsx:repeat> tag syntax iterates over the bean properties row by row, beginning with the start row.

Implicit and explicit indexing

Examples 1, 2, and 3 show how to use the <tsx:repeat> tag. The examples produce the same output if all indexed properties have 300 or fewer elements. If there are more than 300 elements, Examples 1 and 2 display all elements, while Example 3 shows only the first 300 elements.

Example 1 shows *implicit indexing* with the default start and default end index. The bean with the smallest number of indexed properties restricts the number of times the loop repeats.

```
<table>
<tsx:repeat>
  <tr><td><tsx:getProperty name="serviceLocationsQuery" property="city" />
  </tr></td>
  <tr><td><tsx:getProperty name="serviceLocationsQuery" property="address" />
  </tr></td>
  <tr><td><tsx:getProperty name="serviceLocationsQuery" property="telephone" />
  </tr></td>
</tsx:repeat>
</table>
```

Example 2 shows indexing, starting index, and ending index:

```
<table>
<tsx:repeat index=myIndex start=0 end=2147483647>
  <tr><td><tsx:getProperty name="serviceLocationsQuery" property=city(myIndex) />
  </tr></td>
  <tr><td><tsx:getProperty name="serviceLocationsQuery" property=address(myIndex) />
  </tr></td>
  <tr><td><tsx:getProperty name="serviceLocationsQuery" property=telephone(myIndex) />
  </tr></td>
</tsx:repeat>
</table>
```

Example 3 shows *explicit indexing* and ending index with implicit starting index. Although the index attribute is specified, you can still implicitly index the indexed property city because the (myIndex) tag is not required.

```
<table>
<tsx:repeat index=myIndex end=299>
  <tr><td><tsx:getProperty name="serviceLocationsQuery" property="city" /t>
  </tr></td>
  <tr><td><tsx:getProperty name="serviceLocationsQuery" property="address(myIndex)" />
  </tr></td>
  <tr><td><tsx:getProperty name="serviceLocationsQuery" property="telephone(myIndex)" />
  </tr></td>
</tsx:repeat>
</table>
```

Nesting <tsx:repeat> blocks

You can nest <tsx:repeat> blocks. Each block is separately indexed. This capability is useful for interleaving properties on two beans, or properties that have subproperties. In the example, two <tsx:repeat> blocks are nested to display the list of songs on each compact disc in the user's shopping cart.

```
<tsx:repeat index=cdindex>
  <h1><tsx:getProperty name="shoppingCart" property=cds.title /></h1>
  <table>
  <tsx:repeat>
    <tr><td><tsx:getProperty name="shoppingCart" property=cds(cdindex).playlist />
    </td></tr>
  </tsx:repeat>
  </table>
</tsx:repeat>
```

JspBatchCompiler tool

As an IBM enhancement to JavaServer Pages support, IBM WebSphere Application Server provides a batch JSP compiler. Use this function to batch compile your JSP files and thereby enable faster responses to the initial client requests for the JSP files on your production Web server.

Batch compiling makes the first request for a JSP file much faster because the JSP file is translated and compiled into a servlet. Batch compiling is also useful as a fast way to resynchronize all of the JSP files for an application.

To use the JSP batch compiler for JSP files, enter the following command on a single line at an operating system command prompt:

```
JspBatchCompiler -enterpriseapp.name <name>
[ -webmodule.name <name>]
[ -cell.name <name>]
[ -node.name <name>]
[ -server.name <name>]
[ -filename <jsp name>]
[ -keepgenerated <true|false>]
[ -verbose <true|false>]
[ -deprecation <true|false>]
```

If the names specified for these arguments are comprised of two or more words separated by spaces, you must add quotation marks around the names.

where:

- **enterpriseapp.name**
Represents the name of the enterprise application you want to compile.
- **webmodule.name**

Represents the name of the specific Web module that you want to compile. If this argument is not set, all Web modules in the enterprise application are compiled.

- **cell.name**

Represents the name of the cell in which the application is deployed. The default is `BaseApplicationServerCell`.

- **node.name**

Represents the name of the node in which the application is deployed. The default is `DefaultNode`.

- **server.name**

Represents the name of the server in which the application is deployed. The default is `server1`.

- **filename**

Represents the name of a single JSP file that you want to compile. If this argument is not set, all files in the Web module are compiled. Alternatively, if *filename* is set to the name of a directory, only the JSP files in that directory are compiled.

- **keepgenerated**

Represents the option to save or erase the generated files.

If set to yes, WebSphere Application Server saves the generated `.java` files used for compilation on your server. By default, this argument is set to no and the `.java` files are erased after the class files have compiled.

- **verbose**

Indicates the compiler should generate verbose output while compiling the generated sources.

- **deprecation**

Indicates the compiler should generate deprecation warnings while compiling the generated sources.

Bean Scripting Framework

The Bean Scripting Framework (BSF) enables you to use scripting language functions in your Java server-side applications. This framework also extends scripting languages so that you can use existing Java classes and Java beans in the JavaScript language. Support in the JSP Engine for the Bean Scripting Framework is deprecated with WebSphere Application Server 5.1.

With BSF, you can write scripts that create, manipulate and access values from Java objects, or you can write Java programs that evaluate and access results from scripts.

WebSphere Application Server provides the Bean Scripting Framework, which consists of a BSF manager, a BSF engine, and a scripting engine.

BSF provides an access mechanism to Java objects for the scripting languages it supports, so that both the scripting language and the Java code can access code exclusive functions. The access mechanism is implemented through a registry of objects maintained by BSF.

BSF in WebSphere Application Server supports the Rhino ECMAScript.

The "Resources for Learning" article provides external BSF links that document future supported languages.

Developing Web applications

Design a Web application and the components that it needs.

For general Web application design information, see "Resources for learning."

There are two basic approaches to selecting tools for developing Web applications:

- You can use one of the available integrated development environments (IDEs). IDE tools automatically generate significant parts of the servlet and JavaServer Pages (JSP) code, and Hypertext Markup Language (HTML) files. They also contain integrated tools for packaging and testing the Web application components. The IBM WebSphere Application Developer product is the recommended IDE. For more information, see the documentation for that product.
- If you decide to develop Web components without an IDE, you need at least an ASCII text editor. You can also use tools available in the Java Software Development Kit (SDK) and in this product to assemble, test, and deploy the Web application components.

The following steps support the second approach, development without an IDE.

1. If necessary, migrate any pre-existing code to the required version of the servlet and JSP specification.
2. Write and compile the components of the Web application. To access classes that were extended, compile your code using the `-classpath` option on the `javac` compiler. This option allows you to reference the `j2ee.jar` file in the product `<install_root>\lib` directory.

For example, to compile a servlet running on the Windows NT version of WebSphere Application Server, specify:

```
javac -classpath D:\Program Files\WebSphere\AppServer\lib\j2ee.jar MyServlet.java
```

To compile that same servlet on the Windows NT version of WebSphere Network Deployment, specify:

```
javac -classpath D:\Program Files\WebSphere\DeploymentManager\lib\j2ee.jar MyServlet.java
```

3. **(Optional)** Disable JavaServer Pages (JSP) runtime compilation, if necessary.

Assemble the application components in one or more Web modules.

Disabling JavaServer Pages run-time compilation

By default, the JavaServer Pages (JSP) engine translates a requested JSP file, compiles the `.java` file, and loads the compiled servlet into the run-time environment. In previous releases of WebSphere Application Server, if a `.class` file did not exist, the JSP engine always translated and compiled the JSP file. You had to turn off the Web applications reload capability to prevent additional translations and recompiles of the file.

With Version 5.0.1 of WebSphere Application Server, you can now change the JSP engine default behavior by indicating a JSP file should never be translated or compiled at run time, even when a `.class` file does not exist.

If run-time compilation is disabled, you must precompile the JSP files, which provides the following advantages:

- Reduces compilation related disk operations.
- Minimizes disk storage requirements necessary for handling temporary `.java` and `.class` files generated during a run-time compilation.
- Forces you to verify that a JSP file compiled successfully before deploying and installing the application in WebSphere Application Server.

You can disable run-time JSP file compilation on a global or an individual Web application basis:

- To disable the translation and compilation of JSP files for all Web applications, set the Web container Custom property `disableJspRuntimeCompilation` to `true`.

Set this property through the Web container Custom properties panel in the administrative console. To view this administrative console page, click:

```
Servers > Application Servers > server_name > Web Container >
  Custom Properties > property_name
```

Valid values for this setting are `true` or `false`. If this property is set to `true`, then translation and compilation of the JSP files is disabled at run time for all Web applications.

- To disable the translation and compilation of JSP files for a specific Web application, set the JSP engine initialization parameter `disableJspRuntimeCompilation` to `true`. This setting, if enabled, determines the run-time behavior of the JSP engine and overrides the Web container custom property setting.

Set this parameter through the JavaServer Pages attribute assembly settings panel in the Chapter 29, “Assembling applications with the Assembly Toolkit,” on page 1391.

Web Modules > *component_instance* > **Assembly Property Extensions**

Valid values for this setting are `true` or `false`. If this parameter is set to `true`, then, for that specific Web application, translation and compilation of the JSP files is disabled at run time, and the JSP engine only loads precompiled files.

- If neither the Web container custom property nor the JSP attribute assembly parameter is set, the first request for a JSP file results in the translation and compilation of the JSP file when the `.class` file does not exist. Subsequent requests for the file also result in compilations and translations, but only if the following conditions are met:
 - Compilations and translations are required.
 - Reloading is enabled for the Web module.
 - Reload interval is exceeded.

If you disable run-time compilation and a request arrives for a JSP file that does not have a matching `.class` file, the JSP engine returns HTTP error 501 (Not implemented) to the browser. If the JSP file does not exist, the JSP engine returns HTTP error 404 (File not found) to the browser. In both cases, an exception is written to the System Out (SYSOUT) and First Failure Data Capture (FFDC) logs. If a JSP file has a matching `.class` file but that file is out of date, the JSP engine still loads the `.class` file into memory.

Perform the following steps to determine whether the `disableJspRuntimeCompilation` option is enabled in WebSphere Application Server:

1. Enable the Diagnostic Trace Service and set the trace specification to `com.ibm.ws.webcontainer.jsp.servlet.*=all=enabled`.
2. Request a JSP file.
3. Locate the string, `disableJspRuntimeCompilation:true`, in the `trace.log` file.
4. Ensure the `jspUri`: entry matches the requested JSP file.

If both the `disableJspRuntimeCompilation:true` string and the matching `jspUri`: entry appear in the trace, the `disableJspRuntimeCompilation` setting is enabled for the Web application.

Example: Converting JavaScript source to the Bean Scripting Framework

JavaScript code is one of the most popular languages of Web developers. This language supports the following base objects, plus additional objects from the Document Object Model:

- array
- date
- math
- number
- string

Server-side JavaScript code supports the same base objects, and additional objects that support user access to databases, file systems and e-mail systems.

Like client-side JavaScript code, server-side JavaScript code is also platform, browser, and language independent.

You can convert server-side JavaScript applications to the Bean Scripting Framework. This article describes how to perform this conversion.

Server-side JavaScript source code

Suppose you have the following server-side JavaScript application:

```
<html>
<head>
<title>Hello World server-side JavaScript example</title>
</head>
<body>
<br><br>
</body>
</html>

<server>
function writePage()
  write("<center><font size='6'>Hello World</font></center>");
</server>
```

Converting server-side JavaScript source code to the Bean Scripting Framework (BSF)

Make the following changes to the JavaScript source code to enable BSF:

```
<%@ page language="javascript" %>
<html>
<head>
<title>Hello World server-side BSF/JavaScript example</title>
</head>
<body>
<br><br>
</body>
</html>

<%
  out.println("<center><font size='6'>Hello World</font></center>");
%>
```

Review the other BSF reference articles for deployment information and additional programming examples.

Scenario: Creating a Bean Scripting Framework application

Scenario description

Programming skills in JavaScript code are more prevalent than programming skills using JavaServer Pages (JSP) tags. Using the Bean Scripting Framework, JavaScript programmers can gradually introduce JSP tags in their JavaScript applications without completely rewriting the source code. The BSF method not only reduces the potential of programming errors, but also provides a painless way to learn a new technology.

The following scenario illustrates how to implement a BSF application using JavaScript within JSP tags.

Developing the BSF application

At ABC elementary school, John Doe teaches third grade mathematics. He wants to help his students memorize their multiplication tables, and thinks a small Web-based quiz could help meet his objective. However, John Doe only knows JavaScript.

Using the Bean Scripting Framework to help leverage his JavaScript skills, John Doe creates two JSP files, `multiplication_test.jsp` and `multiplication_scoring.jsp`.

In the `multiplication_test.jsp` file, John Doe uses both client-side and server-side JavaScript code to generate a test of 100 random multiplication questions, displayed using a three minute timer. He then writes the `multiplication_scoring.jsp` file to read the data submitted by the `multiplication_test.jsp` file and to generate the scoring results.

John Doe creates the following two files:

multiplication_test.jsp:

```
<html>
<head>
<title>Multiplication Practice Test</title>
<script language="javascript">
var countMin=3;
var countSec=0;
function updateDisplay (min, sec) {
    var disp;
    if (min <= 9) disp = " 0";
    else disp = " ";
    disp += (min + ":");
    if (sec <= 9) disp += ("0" + sec);
    else disp += sec;
    return(disp);
}
function countDown() {
    countSec--;
    if (countSec == -1) {
        countSec = 59;
        countMin--;
    }
    document.multtest.counter.value = updateDisplay(countMin, countSec);
    if((countMin == 0) &&(countSec == 0)) document.multtest.submit();
    else var down = setTimeout("countDown();", 1000);
}
</script>
</head>
<body bgcolor="#ffffff" onLoad="countDown();">
<%@ page language="javascript" %>
<h1>Three Minute Multiplication Drill</h1>
<hr>
<h2>Remember: this is an opportunity to excel!</h2>
<p>
<form method="POST" name="multtest" action="multiplication_scoring.jsp">
<div align="center">
<table>
<tr>
<td>
<h3>Time left:
<input type="text" name="counter" size="9" value="03:00" readonly>
</h3>
</td>
<td>
<input type="submit" value="Submit for scoring!">
</td>
</tr>
</table>
<table border="1">
<%
var newrow = 0;
var q_num = 0;
function addQuestion(num1, num2) {
    if (newrow == 0) out.println("<tr>");
    out.println("<td>");
    out.println(num1 + " x " + num2 + " = ");
    out.println("</td><td>");
    out.print("<input name=\"\" + q_num + \"|\" + num1 + \":\" + num2 + \"\");
    out.println("type=\"text\" size=\"10\">");
    out.println("</td>");

```

```

    if (newrow == 3) {
        out.println("</tr>");
        newrow = 0;
    }
    else newrow++;
    q_num++;
}
for (var i = 0; i < 100; i++) {
    var rand1 = Math.ceil(Math.random() * 12);
    var rand2 = Math.ceil(Math.random() * 12);
    addQuestion(rand1, rand2);
}
%>
</table>
</div>
</form>
</body>
</html>

```

multiplication_scoring.jsp:

```

<html>
<head>
<title>Multiplication Practice Test Results</title>
</head>
<body bgcolor="#ffffff">
<%@ page language="javascript" %>
<h1>Multiplication Drill Score</h1>
<hr>
<div align="center">
<table border="1">
<tr><th>Problem</th><th>Correct Answer</th><th>Your Answer</th></tr>
<%
var total_score = 0;
function score (current, pos1, pos2) {
    var multiplier = current.substring(pos1 + 1, pos2);
    var multiplicand = current.substring(pos2 + 1, current.length());
    var your_product = request.getParameterValues(current)[0];
    var true_product = multiplier * multiplicand;
    out.println("<tr>");
    out.println("<td>" + multiplier + " x " + multiplicand + " = </td>");
    out.println("<td>" + true_product + "</td>");
    if (your_product == true_product) {
        total_score++;
        out.print("<td bgcolor=\"\#00ff00\">");
    }
    else {
        out.print("<td bgcolor=\"\#ff0000\">");
    }
    out.println(your_product + "</td>");
    out.println("</tr>");
}
var equations = request.getParameterNames();
while(equations.hasMoreElements()) {
    var currElt = equations.nextElement();
    var splitPos1 = currElt.indexOf("|");
    var splitPos2 = currElt.indexOf(":");
    if (splitPos1 >=0 && splitPos2 >= 0) score(currElt, splitPos1, splitPos2);
}
%>
</table>
<h2>Total Score: <%= total_score %></h2>
<h3><a href="multiplication_test.jsp">Try again?</a></h3>
</div>
</body>
</html>

```

Follow these steps to see how John Doe uses BSF to implement JavaScript in a JSP application:

1. Give your files a .jsp extension.
2. Use server-side JavaScript code in your application.

The `multiplication_test.jsp` file incorporates both client-side and server-side JavaScript. Server-side JavaScript is similar to client-side JavaScript; the primary difference consists of using a different set of objects. Whereas client-side Javascript programmers invoke `document` and `window` objects, server-side JavaScript programmers, using the Bean Scripting Framework, invoke a set of objects provided by the JSP technology. Also, client-side scripts are enclosed in `<script>` tags, but server-side scripts use JSP scriptlet and expression tags.

3. Examine the following blocks of code:

```
<script language="javascript">
var countMin=3;
var countSec=0;
function updateDisplay (min, sec) {
    var disp;
    if (min <= 9) disp = " 0";
    else disp = " ";
    disp += (min + ":");
    if (sec <= 9) disp += ("0" + sec);
    else disp += sec;
    return(disp);
}
function countDown() {
    countSec--;
    if (countSec == -1) {
        countSec = 59;
        countMin--;
    }
    document.mulptest.counter.value = updateDisplay(countMin, countSec);
    if((countMin == 0) && (countSec == 0)) document.mulptest.submit();
    else var down = setTimeout("countDown();", 1000);
}
</script>
....
<body bgcolor="#ffffff" onLoad="countDown();">
...
<form method="POST" name="multtest" action="multiplication_scoring.jsp">
...
<input type="text" name="counter" size="9" value="03:00" readonly>
...
```

The JavaScript code contained in the `<script>` block implements a timer set within the `<input>` field named `counter`. The `onLoad` event handler in the `<body>` tag causes the browser to load and execute the code when the the page is loaded.

The `document.mulptest.submit()` statement causes the form named `multtest` to be submitted when the timer expires.

4. Identify the code to the BSF function.

The following code example, from the `multiplication_test.jsp` file, displays the use of a JSP directive. This directive tells the WebSphere Application Server BSF function that this file is using the JavaScript language, and that the JavaScript code is enclosed by the `<% ... %>` scriptlet tags. The out implicit JSP object in this code example, creates the body section of a table from 100 randomly generated questions.

```
...
<%@ page language="javascript" %>
...
<%
var newrow = 0;
var q_num = 0;

function addQuestion(num1, num2) {
    if (newrow == 0) out.println("<tr>");
```

```

        out.println("<td>");
        out.println(num1 + " x " + num2 + " = ");
        out.println("</td><td>");
        out.print("<input name=\"" + q_num + "\" |" + num1 + ":" + num2 + "\" ");
        out.println("type=\"text\" size=\"10\">");
        out.println("</td>");

        if (newrow == 3) {
            out.println("</tr>");
            newrow = 0;
        }
        else newrow++;

        q_num++;
    }

    for (var i = 0; i < 100; i++) {
        var rand1 = Math.ceil(Math.random() * 12);
        var rand2 = Math.ceil(Math.random() * 12);

        addQuestion(rand1, rand2);
    }

    %>
    ...

```

5. Read the results.

To score the results of the practice drill, John Doe uses the request implicit JSP object in the `multiplication_scoring.jsp` file to obtain the POST data created within the `<form>` tags in the `multiplication_test.jsp` file.

The `multiplication_scoring.jsp` file uses the POST data to build an output file containing the original question, the student's answer, and the correct answer, and then prints the text in a table format using the `out` implicit object.

The following code example from the `multiplication_scoring.jsp` file illustrates the use of the `request` and `out` JSP objects:

```

...
<%@ page language="javascript" %>
...
<%
var total_score = 0;
function score (current, pos1, pos2) {
    var multiplier = current.substring(pos1 + 1, pos2);
    var multiplicand = current.substring(pos2 + 1, current.length());
    var your_product = request.getParameterValues(current)[0];
    var true_product = multiplier * multiplicand;
    out.println("<tr>");
    out.println("<td>" + multiplier + " x " + multiplicand + " = </td>");
    out.println("<td>" + true_product + "</td>");
    if (your_product == true_product) {
        total_score++;
        out.print("<td bgcolor=\"\#00ff00\">");
    }
    else {
        out.print("<td bgcolor=\"\#ff0000\">");
    }
    out.println(your_product + "</td>");
    out.println("</tr>");
}
var equations = request.getParameterNames();
while(equations.hasMoreElements()) {
    var currElt = equations.nextElement();
    var splitPos1 = currElt.indexOf("|");
    var splitPos2 = currElt.indexOf(":");
    if (splitPos1 >=0 && splitPos2 >= 0) score(currElt, splitPos1, splitPos2);
}

```



```

%>
...
<h2>Total Score: <%= total_score %></h2>
...

```

Note: Although using separate scriptlet blocks of code for different portions of a conditional expression is common in JSP files implemented in Java, it is invalid for JSP files implemented using JavaScript through the Bean Scripting Framework. The JavaScript code must be entirely contained within the scriptlet tags.

The following code example illustrates invalid usage:

```

<% if (pass == 0) %>
    <i>pass is true</i>
<% else %>
    <i>pass is not true</i>

```

Deploying the BSF application

You assemble and deploy BSF applications in the same manner as JSP applications. Review the Chapter 29, “Assembling applications with the Assembly Toolkit,” on page 1391 article for more information.

Deploy the BSF code examples in WebSphere Application Server to view this applications processing and output. Use the following quick steps to deploy the application.

The intent of these “quick steps” is to provide you with instant application output. However, the supported method for deployment is the same as for standard JSP files.

1. Use the DefaultApplication to add your BSF files.

Copy your .jsp files to the DefaultApplication directory:

```

<app server install directory>/installedApps/<node
name>/DefaultApplication.ear/DefaultApplication.war

```

2. Start the application server.
3. Open a browser and request your BSF application.

Use the following URL to request your application:

```

http://hostName:9080/<jspFileName>.jsp

```

Example: Bean Scripting Framework code example

The following code examples show how to implement JavaScript using the Bean Scripting Framework (BSF).

For a quick demonstration of the BSF function, copy these code examples into 2 separate files, and deploy them in WebSphere Application Server using the instructions in the BSF scenario article.

Multiplication practice test

```

<html>
<head>
<title>Multiplication Practice Test</title>
<!--
This file and its companion, multiplication_score.jsp, illustrate the
use of ECMAScript within the BSF framework. The task is a simple
timed math quiz, which is 3 minutes in duration. When the quiz ends,
the score is computed and displayed.
Users are then asked if they wish to try
the quiz again.
-->
<!--

```

This code fragment displays and updates the quiz countdown in client side JavaScript code.

```
-->
<script language="javascript">
var countMin=3;
var countSec=0;

// This code computes the current countdown time.
function updateDisplay (min, sec) {
    var disp;

    if (min <= 9) disp = " 0";
    else disp = " ";

    disp += (min + ":");

    if (sec <= 9) disp += ("0" + sec);
    else disp += sec;

    return(disp);
}

//This code fragment displays the current countdown time in the user's
//browser window, and submits the results for scoring when the countdown
//ends.

function countDown() {
    countSec--;
    if (countSec == -1) {
        countSec = 59;
        countMin--;
    }
    document.mutttest.counter.value = updateDisplay(countMin, countSec);
    if((countMin == 0) && (countSec == 0)) document.mutttest.submit();
    else var down = setTimeout("countDown();", 1000);
}

</script>
</head>
<body bgcolor="#ffffff" onLoad="countDown();">

<!--
The body of the quiz runs as JavaServer Pages (JSP) code using BSF.
The code outputs the problems in table format using the POST method
and invokes the scoring module when the user chooses to end the quiz
or when the countdown ends.
-->
<%@ page language="javascript" %>

<h1>Three Minute Multiplication Drill</h1>
<hr>

<h2>Remember: this is an opportunity to excel!</h2>
<p>

<form method="POST" name="multttest" action="multiplication_scoring.jsp">
<div align="center">
<table>
<tr>
<td>
<h3>Time left:
<input type="text" name="counter" size="9" value="03:00" readonly>
</h3>
</td>
<td>
<input type="submit" value="Submit for scoring!">
</td>

```

```

</tr>
</table>
<table border="1">
<%
var newrow = 0;
var q_num = 0;

// This code generates a new random multiplication problem up to the number
//twelve, and enters it into the table of problems.

function addQuestion(num1, num2) {
    if (newrow == 0) out.println("<tr>");

    out.println("<td>");
    out.println(num1 + " x " + num2 + " = ");
    out.println("</td><td>");
    out.print("<input name=\"" + q_num + "\"" + num1 + ":" + num2 + "\" ");
    out.println("type=\"text\" size=\"10\">");
    out.println("</td>");

    if (newrow == 3) {
        out.println("</tr>");
        newrow = 0;
    }
    else newrow++;

    q_num++;
}

//This code obtains two random operands and formats 100 quiz problems.

for (var i = 0; i < 100; i++) {
    var rand1 = Math.ceil(Math.random() * 12);
    var rand2 = Math.ceil(Math.random() * 12);

    addQuestion(rand1, rand2);
}

%>
</table>
</div>
</form>

</body>
</html>

```

Multiplication practice test results

```

<html>
<head>
<title>Multiplication Practice Test Results</title>
</head>
<body bgcolor="#ffffff">

<!--
This JSP code is invoked when the user submits a math quiz for scoring,
or when the quiz countdown expires. The JSP code tabulates the problem list,
the correct answer, the user's answer, and scores the test. It then offers
the user an opportunity to try the quiz again.
-->
<%@ page language="javascript" %>

<h1>Multiplication Drill Score</h1>
<hr>

<div align="center">

```

```

<table border="1">
<tr><th>Problem</th><th>Correct Answer</th><th>Your Answer</th></tr>
<%
var total_score = 0;

// This code parses the submitted form, extracts the a problem generated by the
// multiplication_test.jsp file, outputs it, computes the correct answer,
// and displays this information and the user answer. The code scores
// the quiz using a running sum of correct answers.

function score (current, pos1, pos2) {
    var multiplier = current.substring(pos1 + 1, pos2);
    var multiplicand = current.substring(pos2 + 1, current.length());
    var your_product = request.getParameterValues(current)[0];
    var true_product = multiplier * multiplicand;

    out.println("<tr>");
    out.println("<td>" + multiplier + " x " + multiplicand + " = </td>");
    out.println("<td>" + true_product + "</td>");

    if (your_product == true_product) {
        total_score++;
        out.print("<td bgcolor=\"\#00ff00\">");
    }
    else {
        out.print("<td bgcolor=\"\#ff0000\">");
    }
    out.println(your_product + "</td>");
    out.println("</tr>");
}

// This is the main body of the scoring application. It parses the posted quiz,
// and calls the score() function to score remaining problems.

var equations = request.getParameterNames();
while(equations.hasMoreElements()) {
    var currElt = equations.nextElement();
    var splitPos1 = currElt.indexOf("|");
    var splitPos2 = currElt.indexOf(":");

    if (splitPos1 >=0 && splitPos2 >= 0) score(currElt, splitPos1, splitPos2);
}

%>
</table>

<h2>Total Score: <%= total_score %></h2>
<h3><a href="/multiplication_test.jsp">Try again?</a></h3>
</div>

</body>
</html>

```

Web modules

A Web module represents a Web application. A Web module is created by assembling servlets, JavaServer Pages (JSP) files, and static content such as HyperText Markup Language (HTML) pages into a single deployable unit. Web modules are stored in Web archive (WAR) files, which are standard Java archive files.

A Web module contains:

- One or more servlets, JSP files, and HTML files.
- A deployment descriptor, stored in an Extensible Markup Language (XML) file.

The file, named `web.xml`, declares the contents of the module. It contains information about the structure and external dependencies of Web components in the module and describes how the components are used at run time.

You can create Web modules as stand-alone applications, or you can combine Web modules with other modules to create J2EE applications. You install and run a Web module in the Web container of an application server.

Assembling Web applications

Assemble a Web module to contain servlets, JavaServer page (JSP) files, and related code artifacts. (Group enterprise beans, client code, and resource adapter code in separate modules). After assembling a Web module, you can install it as a stand-alone application or combine it with other modules into an enterprise application.

Use the Assembly Toolkit to assemble a Web module in any of the following ways:

- Import an existing Web module (WAR file).
- Create a new Web module.
- Copy code artifacts (such as servlets) from one Web module into a new Web module.

Although you can input various properties for Web archives, available properties are specific to the Servlet, JSP, and J2EE specification level.

1. Start the Assembly Toolkit.
2. Optional: Open the J2EE perspective to work with J2EE projects. Click **Window > Open Perspective > Other > J2EE**.
3. Optional: Open the J2EE Hierarchy view. Click **Window > Show View > J2EE Hierarchy**. Other helpful views include the Project Navigator view (**Window > Show View > Other > J2EE > Project Navigator**) and the Navigator view (**Window > Show View > Navigator**).
4. **5.1+** Optional: To use application assembly services of WebSphere Business Integration (WBI) Server Foundation such as Internationalization, ActivitySession or Application profiling, select the **Integration Server** target server.
 - a. Click **Project > Properties > J2EE**.
 - b. For **Target server**, select **Integration Server**. For this release, **Integration Server v5.1**.
 - c. Click **OK**.

You can set values for WBI Server Foundation services on the **Extended Services** tab of a deployment descriptor editor open on the module.

5. Migrate WAR files created with the Application Assembly Tool (AAT) or a different tool to the Assembly Toolkit. To migrate files, import your WAR files to the Assembly Toolkit.
6. Create a new Web module.
7. Copy code artifacts (such as servlets) from one Web module into a new Web module.
8. Verify the contents of the new Web module in either of the following ways:
 - In the J2EE Hierarchy view, expand **Web Modules** and view the new module.
 - Click **Window > Show View > Navigator** to see the associated files for the Web module in a Navigator view.

Context parameters

A servlet context defines a server's view of the Web application within which the servlet is running. The context also allows a servlet to access resources available to it.

Using the context, a servlet can log events, obtain URL references to resources, and set and store attributes that other servlets in the context can use. These properties declare a Web application's parameters for its context. They convey setup information, such as a webmaster's e-mail address or the name of a system that holds critical data.

Security constraints

Security constraints determine how Web content is to be protected.

These properties associate security constraints with one or more Web resource collections. A constraint consists of a Web resource collection, an authorization constraint and a user data constraint.

- A Web resource collection is a set of resources (URL patterns) and HTTP methods on those resources. All requests that contain a request path that matches the URL pattern described in the Web resource collection is subject to the constraint. If no HTTP methods are specified, then the security constraint applies to all HTTP methods.
- An authorization constraint is a set of roles that users must be granted in order to access the resources described by the Web resource collection. If a user who requests access to a specified URI is not granted at least one of the roles specified in the authorization constraint, the user is denied access to that resource.
- A user data constraint indicates that the transport layer of the client or server communications process must satisfy the requirement of either guaranteeing content integrity (preventing tampering in transit) or guaranteeing confidentiality (preventing reading while in transit).

Servlet mappings

A servlet mapping is a correspondence between a client request and a servlet.

Servlet containers use URL paths to map client requests to servlets, and follow the URL path-mapping rules as specified in the Java Servlet specification. The container uses the URI from the request, minus the context path, as the path to map to a servlet. The container chooses the longest matching available context path from the list of Web applications that it hosts.

Invoker attributes

Invoker attributes are used by the servlet that implements the invocation behavior.

Error pages

Error page locations allow a servlet to find and serve a URI to a client based on a specified error status code or exception type.

These properties are used if the error handler is another servlet or JSP file. The properties specify a mapping between an error code or exception type and the path of a resource in the Web application. The container examines the list in the order that it is defined, and attempts to match the error condition by status code or by exception class. On the first successful match of the error condition, the container serves back the resource defined in the Location property.

File serving

File serving allows a Web application to serve static file types, such as HTML. File-serving attributes are used by the servlet that implements file-serving behavior.

Initialization parameters

Initialization parameters are sent to a servlet in its `HttpConfig` object when the servlet is first started.

Servlet caching

Dynamic caching can be used to improve the performance of servlet and JavaServer Pages (JSP) files by serving requests from an in-memory cache. Cache entries contain the servlet's output, results of the servlet's execution, and metadata.

Web components

A web component is a servlet, Java Server Page (JSP), or HTML file. One or more web components make up a web module.

Web property extensions

Web property extensions are IBM extensions to the standard deployment descriptors for Web applications. These extensions include mime filtering and servlet caching.

Web resource collections

A Web resource collection defines a set of URL patterns (resources) and HTTP methods belonging to the resource.

HTTP methods handle HTTP-based requests, such as GET, POST, PUT, and DELETE. A URL pattern is a partial Uniform Resource Locator that acts as a template for matching the pattern with existing full URLs in an attempt to find a valid file.

Welcome files

A Welcome file is an entry point file (for example, `index.html`) for a group of related HTML files.

Welcome files are located by using a group of partial URIs. The Web container uses the partial URIs to find a valid file when the initial URI is not found.

Troubleshooting tips for Web application deployment

Deployment of a Web application is successful if you can access the application by typing a Uniform Resource Locator (URL) in a browser, or if you can access the application by following a link.

If you cannot access your application, follow these steps to eliminate some common errors that can occur during migration or deployment.

Web module does not run in WebSphere Application Server Version 5.

Symptom	Your Web module does not run when you migrate it to Version 5
Problem	In Version 4.x, the classpath setting that affected visibility was <i>Module Visibility Mode</i> . In Version 5, you must use class loader policies to set visibility.
Recommended response	Reassemble an existing module, or change the visibility settings in the class loader policies. in the class loader policies.
	See article Migration of module visibility modes from Version 4.x for more information and examples.

Welcome page is not visible.

Symptom	You cannot access an application with a Web path of: <code>/webapp/myapp</code>
Problem	The default welcome page for a Web application is assumed to be <i>index.html</i> . You cannot access the default page of the <i>myapp</i> application unless it is named <i>index.html</i> .

Recommended response To identify a different welcome page, modify the properties of the Web module during assembly.

HTML files are not found.

Symptom Your Web application ran successfully on prior versions, but now you encounter errors that the welcome page (typically *index.html*), or referenced HTML files are not found:
Error 404: File not found: Banner.html
Error 404: File not found: HomeContent.html

Problem For security and consistency reasons, Web application URLs are now case-sensitive on all operating systems.

Suppose the content of the index page is as follows:

```
<!DOCTYPE HTML PUBLIC "-//W3C//DTD HTML 5.0 Frameset//EN">
<HTML>
<TITLE>
Insurance Home Page
</TITLE>
  <frameset rows="18,80">
    <frame src="Banner.html" name="BannerFrame" SCROLLING=NO>
    <frame src="HomeContent.html" name="HomeContentFrame">
  </frameset>
</HTML>
```

However the actual file names in the \WebSphere\AppServer\installedApps\... directory where the application is deployed are:

```
banner.html
homecontent.html
```

Recommended response To correct this problem, modify the *index.html* file to change the names *Banner.html* and *HomeContent.html* to *banner.html* and *homecontent.html* to match the names of the files in the deployed application.

For current information available from IBM Support on known problems and their resolution, see the IBM Support page.

IBM Support has documents that can save you time gathering information needed to resolve this problem. Before opening a PMR, see the IBM Support page.

Modifying the default Web container configuration

The Web container is created initially with default properties values suitable for simple Web applications. However, these values might not be appropriate for more complex Web applications.

Your application is considered complex if it requires any of the following features:

- virtual host
- servlet caching
- special client request loads
- persistent HTTP session support
- special HTTP transport settings
- transaction class mappings

Modify the following properties if you have a complex application:

1. If your Web application requires a virtual host, other than the default_host, or requires servlet caching, modify the Web container **General Properties**.
2. If your application handles special client request loads, modify the Web Container **Additional Properties > Thread Pool** setting.

3. If your application requires persistent HTTP session support, modify the Web Container **Additional Properties > Session Management** setting.
4. If your application requires one of the following HTTP transport settings:
 - Unique hostname and port for client access
 - SSL enablementmodify the Web Container **Additional Properties > HTTP transports** setting.
5. If your application requires global settings for internal servlets for WAR files packaged by third-party tools, modify the Web Container **Additional Properties > Custom Properties** setting.
6. If your application uses transaction class mappings to classify workload, modify the Web Container **Additional Properties > Advanced Settings**.

Web container

A Web container handles requests for servlets, JavaServer Pages (JSP) files, and other types of files that include server-side code. The Web container creates servlet instances, loads and unloads servlets, creates and manages request and response objects, and performs other servlet management tasks.

The Web server plug-ins, provided by the WebSphere Application Server, help supported Web servers pass servlet requests to Web containers.

Web container settings

Use this page to configure the web container settings.

To view this administrative console page, click **Servers > Application Servers > *server_instance* > Web container**.

Configuration - General Properties

Default virtual host

Specifies a virtual host that enables a single host machine to resemble multiple host machines. Resources associated with one virtual host cannot share data with resources associated with another virtual host, even if the virtual hosts share the same physical machine.

Select a virtual host option:

Default Host

The product provides a default virtual host with some common aliases, such as the machine IP address, short host name, and fully qualified host name. The alias comprises the first part of the path for accessing a resource such as a servlet. For example, it is `localhost:9080` in the request `http://localhost:9080/myServlet`.

Admin Host

This is another name for the application server; also known as *server1* in the base installation. This process supports the use of the administrative console.

Servlet caching

Specifies that if a servlet is invoked once and it generates output to be cached, a cache entry is created containing not only the output, but also side effects of the invocation. These side effects can include calls to other servlets or Java Server Pages (JSP) files, as well as metadata about the entry, including timeout and entry priority information.

Enable servlet caching

Check this box to enable servlet caching.

Web Module Deployment settings

Use this page to configure an instance of Web module deployment.

To view this administrative console page, click **Applications > Enterprise Application > application_instance > Web Modules > Web Module_instance**.

URI

Specifies a URI that, when resolved relative to the application URL, specifies the location of the module archive contents on a file system. The URI must match the ModuleRef URI in the deployment descriptor of an application if the module was packaged as part of a deployed application or enterprise archive (EAR) file.

Alternate DD

Specifies the file name for an alternative deployment descriptor file to use instead of the original deployment descriptor file in the module JAR file.

This file is the *post-assembly* version of the deployment descriptor file. You can edit the original deployment descriptor file to resolve dependencies and security information. Specifying the use of the alternative deployment descriptor keeps the original deployment descriptor file intact.

The value of the *Alternate DD* property must be the full path name of the deployment descriptor file, relative to the module root directory. By convention, the file is in the ALT-INF directory. If this property is not specified, the deployment descriptor file is read from the module JAR file.

Starting weight

Specifies the order in which modules are started. Lower weighted modules are started before higher weighted modules.

Classloader Mode

Specifies whether the class loader should search in the parent class loader or in the application class loader first to load a class. The standard for JDK class loaders and WebSphere class loaders is PARENT_FIRST. By specifying PARENT_LAST, your application can override classes contained in the parent class loader, but this action can potentially result in ClassCastException or LinkageErrors if you have mixed use of overridden classes and non-overridden classes.

The options are PARENT_FIRST and PARENT_LAST. The default is to search in the parent class loader before searching in the application class loader to load a class.

Data type	String
Default	PARENT_FIRST

Web container custom properties

Use this page to configure arbitrary name-value pairs of data, where the name is a property key and the value is a string value that can be used to set internal system configuration properties. Defining a new property enables you to configure a setting beyond that which is available in the administrative console.

To view this administrative console page, click **Servers > Application Servers > server_name > Web Container > Custom Properties**.

Name

Specifies the name (or key) for the property.

Data type	String
------------------	--------

Value

Specifies the value paired with the specified name.

Data type	String
------------------	--------

Description

Provides information about the name-value pair.

Data type String

Global settings for internal servlets

Web Archive (WAR) files packaged using third-party tools cannot specify behavior for the services exposed by the Web container internal servlets. You can globally enable/disable internal servlets for all Web applications at the Web container level by creating name-value pairs such as:

<u>Name</u>	<u>Value</u>
fileServingEnabled	true
directoryBrowsingEnabled	true
serveServletsByClassnameEnabled	true

Settings defined in an assembly tool take precedence over the global settings set through the custom properties at the Web container level.

Web application deployment extensions continue to hold configuration information for the services provided by the internal servlets, and take precedence over the global settings set through the custom properties at the Web container level.

UTF-8 encoded URLs

WebSphere Application Server Version 5.1, introduces support for UTF-8 encoded Uniform Resource Locators (URLs) to support the double-byte characters in URLs. The UTF-8 encoded URL feature is enabled by default. You can prevent the web container from explicitly decoding URLs in UTF-8 and have them use the ISO-8859 standard as per the current HTTP specification by using the following name-value pair:

<u>Name</u>	<u>Value</u>
DecodeUriAsUTF8	false

Web applications: Resources for learning

Use the following links to find relevant supplemental information about Web applications. The information resides on IBM and non-IBM Internet sites, whose sponsors control the technical accuracy of the information.

These links are provided for convenience. Often, the information is not specific to the IBM WebSphere Application Server product, but is useful all or in part for understanding the product. When possible, links are provided to technical papers and Redbooks that supplement the broad coverage of the release documentation with in-depth examinations of particular product areas.

View links to additional information about:

- “Web applications: Resources for learning”
- “Web applications: Resources for learning”
- “Web applications: Resources for learning”

Programming model and decisions

- J2EE BluePrints for Web applications
- Redbook on the design and implementation of Servlets, JSP files, and enterprise beans

Programming instructions and examples

- Redbook on Servlet and JSP file Programming
- Sun's Java™ Tutorial on Servlets
- Introduction to JavaServer Pages - Tutorial
- Bean Scripting Framework description
- Web delivered samples in the Samples Gallery

Programming specifications

- Java 2 Software Development Kit (SDK)
- Servlet 2.3 Specification
- JavaServer Pages 1.2 Specification
- Differences between JavaScript and ECMAScript
- ISO 8859 Specifications

Chapter 3. Managing HTTP sessions

IBM WebSphere Application Server provides a service for managing HTTP sessions: Session Manager. The key activities for session management are summarized below.

Before you begin these steps, make sure you are familiar with the programming model for accessing HTTP session support in the applications following the Servlet 2.3 API.

1. Plan your approach to session management, which could include session tracking, session recovery, and session clustering.
2. Create or modify your own applications to use session support to maintain sessions on behalf of Web applications.
3. Assemble your application.
4. Deploy your application.
5. Ensure the administrator appropriately configures session management in the administrative domain.
6. Adjust configuration settings and perform other tuning activities for optimal use of sessions in your environment.

Sessions

A session is a series of requests to a servlet, originating from the same user at the same browser.

Sessions allow applications running in a Web container to keep track of individual users.

For example, a servlet might use sessions to provide "shopping carts" to online shoppers. Suppose the servlet is designed to record the items each shopper indicates he or she wants to purchase from the Web site. It is important that the servlet be able to associate incoming requests with particular shoppers. Otherwise, the servlet might mistakenly add Shopper_1's choices to the cart of Shopper_2.

A servlet distinguishes users by their unique session IDs. The session ID arrives with each request. If the user's browser is cookie-enabled, the session ID is stored as a cookie. As an alternative, the session ID can be conveyed to the servlet by URL rewriting, in which the session ID is appended to the URL of the servlet or JavaServer Pages (JSP) file from which the user is making requests. For requests over HTTPS or Secure Sockets Layer (SSL), Another alternative is to use SSL information to identify the session.

Migrating HTTP sessions

Note: In Version 5 default write frequency mode is `TIME_BASED_WRITES`, which is different from Version 4.0 and 3.5 default mode of `END_OF_SERVICE`.

Migrating from Version 4.0

No programmatic changes are required to migrate from version 4.0 to version 5.

Migrating from Version 3.5

If you have Version 3.5 applications running in Servlet 2.1 mode, some of the following Version 5 differences might influence how you choose to track and manage sessions.

1. During application development, modify session-related APIs as needed.
Some API changes are required in order to redeploy existing applications on Version 5. These include changes to the `HttpSession` API itself as well as issues associated with moving to support for the

Servlet 2.3 specification. Certain Servlet 2.1 API methods have been deprecated in Servlet 2.3 API . These deprecated APIs still work in Version 5.0, but they may be removed in a future version of the API. Changes are summarized in the following list:

- Replace instances of `getValue()` with `getAttribute()`
- Replace instances of `getValueNames()` with `getAttributeNames()`
- Replace instances of `removeValue()` with `removeAttribute()`
- Replace instances of `putValue()` with `setAttribute()`

2. During application development, modify Web application behavior as needed.

In accordance with the Servlet 2.3 specification, `HttpSession` objects must be scoped within a single Web application context; they may not be shared between contexts. This means that a session can no longer span Web applications. Objects added to a session by a servlet or JSP in one Web application cannot be accessed from another Web application. The same session ID may be shared (because the same cookie is in use), but each Web application will have a unique session associated with the session ID. Version 5 provides a feature that can be used to extend scope of a session to enterprise application.

3. Use administrative tools to configure Session Manager security settings as needed. Relative to session security, the default Session Manager setting for Integrate Security is now false. This is different from the default setting in some earlier releases.
4. Use administrative tools to configure the JSP enabler and application server as needed.

In Version 3.5 of the product, JSP files that contained the usebean tag with scope set to session did not always work properly when session persistence was enabled. Specifically, the JSP writer needed to write a scriptlet to explicitly set the attribute (that is, to call `setAttribute()`) if it was changed as part of JSP processing.

Two new features in Version 5.0 help address this problem:

- You can set `dosetattribute` to true on the JSP InitParameter.
- You can set the Write Contents option to Write all.

The differences between the two solutions are summarized in the following table:

	Applies to	Configured at	Action
<code>dosetattribute</code> set to <i>true</i>	JSP	JSP enabler	Assures that JSP session-scoped beans always call <code>setAttribute()</code>
Write Contents option set to <i>Write all</i>	servlet or JSP	application server	All session data (changed or unchanged) is written to the external location

If session persistence is enabled and a class reload for the Web application occurs, the sessions associated with the Web application are maintained in the persistent store and will be available after the reload.

Developing session management in servlets

This information, combined with the coding example `SessionSample.java`, provides a programming model for implementing sessions in your own servlets.

1. Get the `HttpSession` object.

To obtain a session, use the `getSession()` method of the `javax.servlet.http.HttpServletRequest` object in the Java Servlet 2.3 API.

When you first obtain the `HttpSession` object, the Session Management facility uses one of three ways to establish tracking of the session: cookies, URL rewriting, or Secure Sockets Layer (SSL) information.

Assume the Session Management facility uses cookies. In such a case, the Session Management facility creates a unique session ID and typically sends it back to the browser as a *cookie*. Each

subsequent request from this user (at the same browser) passes the cookie containing the session ID, and the Session Management facility uses this ID to find the user's existing HttpSession object.

In Step 1 of the code sample, the Boolean(create) is set to true so that the HttpSession object is created if it does not already exist. (With the Servlet 2.3 API, the javax.servlet.http.HttpServletRequest.getSession() method with no boolean defaults to true and creates a session if one does not already exist for this user.)

2. Store and retrieve user-defined data in the session.

After a session is established, you can add and retrieve user-defined data to the session. The HttpSession object has methods similar to those in java.util.Dictionary for adding, retrieving, and removing arbitrary Java objects.

In Step 2 of the code sample, the servlet reads an integer object from the HttpSession, increments it, and writes it back. You can use any name to identify values in the HttpSession object. The code sample uses the name sessiontest.counter.

Because the HttpSession object is shared among servlets that the user might access, consider adopting a site-wide naming convention to avoid conflicts.

3. (Optional) Output an HTML response page containing data from the HttpSession object.

4. Provide feedback to the user that an action has taken place during the session. You may want to pass HTML code to the client browser indicating that an action has occurred. For example, in step 3 of the code sample, the servlet generates a Web page that is returned to the user and displays the value of the sessiontest.counter each time the user visits that Web page during the session.

5. (Optional) Notify Listeners. Objects stored in a session that implement the javax.servlet.http.HttpSessionBindingListener interface are notified when the session is preparing to end and become invalidated. This notice enables you to perform post-session processing, including permanently saving the data changes made during the session to a database.

6. End the session. You can end a session:

- Automatically with the Session Management facility if a session is inactive for a specified time. The administrators provide a way to specify the amount of time after which to invalidate a session.
- By coding the servlet to call the invalidate() method on the session object.

Example: SessionSample.java

```
import java.io.*;
import java.util.*;
import javax.servlet.*;
import javax.servlet.http.*;

public class SessionSample extends HttpServlet {
    public void doGet (HttpServletRequest request, HttpServletResponse response)
        throws ServletException, IOException {

        // Step 1: Get the Session object

        boolean create = true;
        HttpSession session = request.getSession(create);

        // Step 2: Get the session data value

        Integer ival = (Integer)
            session.getAttribute ("sessiontest.counter");
        if (ival == null) ival = new Integer (1);
        else ival = new Integer (ival.intValue () + 1);
        session.setAttribute ("sessiontest.counter", ival);

        // Step 3: Output the page

        response.setContentType("text/html");
```

```

    PrintWriter out = response.getWriter();
    out.println("<html>");
    out.println("<head><title>Session Tracking Test</title></head>");
    out.println("<body>");
    out.println("<h1>Session Tracking Test</h1>");
    out.println("You have hit this page " + ival + " times" + "<br>");
    out.println("Your " + request.getHeader("Cookie"));
    out.println("</body></html>");
}
}

```

Assembling so that session data can be shared

In accordance with the Servlet 2.3 API specification, by default the Session Management facility supports session scoping by Web module. Only servlets in the same Web module can access the data associated with a particular session. WebSphere Application Server provides an option that you can use to extend the scope of the session attributes to an enterprise application. Therefore, you can share session attributes across all the Web modules in an enterprise application. This option is provided as an IBM extension.

Restriction: To use this option, you must install all the Web modules in the enterprise application on a given server. You cannot split up Web modules in the enterprise application by servers. For example, with an enterprise application containing two Web modules, you cannot use this option when one Web module is installed on one server and second Web module is installed on a different server. In such split installations, applications might share session attributes across Web modules using distributed sessions, but session data integrity is lost when concurrent access to a session is made in different Web modules. It also severely restricts use of some Session Management features, like TIME_BASED_WRITES. For enterprise applications on which this option is enabled, the Session Management configuration on the Web module inside the enterprise application is ignored. Then Session Management configuration defined on enterprise application is used if Session Management is overwritten at the enterprise application level. Otherwise, the Session Management configuration on the Web container is used.

Servlet API Behavior

Note: If shared HttpSession context is turned on in an enterprise application, HttpSession listeners defined in all the Web modules inside the enterprise application are invoked for session events. The order of listener invocation is not guaranteed.

Do the following to share session data across Web modules in an enterprise application:

1. Start the Assembly Toolkit.
2. In the Assembly Toolkit, right-click the application (EAR file) you want to share and click **Open With > Deployment Descriptor Editor**.
3. In the application deployment descriptor editor of the Assembly Toolkit, select **Shared session context** under **WebSphere Extensions**. Make sure the class definition of attributes put into session are available to all Web modules in the enterprise application. The shared session context does not fully meet the requirements of the Specifications.
4. Save the application (EAR) file. In the Assembly Toolkit, after you close the application deployment descriptor editor, confirm that you want to save changes made to the application.

Session security support

You can integrate HTTP sessions and security in IBM WebSphere Application Server. When security integration is enabled in the Session Management facility and a session is accessed in a protected resource, you can access that session only in protected resources from then on. You cannot mix secured and unsecured resources accessing sessions when security integration is turned on. Security integration in the Session Management facility is not supported in form-based login with SWAM.

Security integration rules for HTTP sessions

Only authenticated users can access sessions created in secured pages and are created under the identity of the authenticated user. Only this authenticated user can access these sessions in other secured pages. To protect these sessions from unauthorized users, you cannot access them from an unsecured page.

Programmatic details and scenarios

IBM WebSphere Application Server maintains the security of individual sessions.

An identity or user name, readable by the `com.ibm.websphere.servlet.session.IBMSession` interface, is associated with a session. An unauthenticated identity is denoted by the user name `anonymous`. IBM WebSphere Application Server includes the `com.ibm.websphere.servlet.session.UnauthorizedSessionRequestException` class, which is used when a session is requested without the necessary credentials.

The Session Management facility uses the WebSphere Application Server security infrastructure to determine the authenticated identity associated with a client HTTP request that either retrieves or creates a session. WebSphere Application Server security determines identity using certificates, LPTA, and other methods.

After obtaining the identity of the current request, the Session Management facility determines whether to return the session requested using a `getSession()` call or not.

The following table lists possible scenarios in which security integration is enabled with outcomes dependent on whether the HTTP request is authenticated and whether a valid session ID and user name was passed to the Session Management facility.

	Unauthenticated HTTP request is used to retrieve a session	HTTP request is authenticated, with an identity of "FRED" used to retrieve a session
No session ID was passed in for this request, or the ID is for a session that is no longer valid	A new session is created. The user name is <code>anonymous</code>	A new session is created. The user name is <code>FRED</code>
A session ID for a valid session is passed in. The current session user name is <code>"anonymous"</code>	The session is returned.	The session is returned. Session Management changes the user name to <code>FRED</code>
A session ID for a valid session is passed in. The current session user name is <code>FRED</code>	The session is not returned. An <code>UnauthorizedSessionRequestException</code> error is thrown*	The session is returned.
A session ID for a valid session is passed in. The current session user name is <code>BOB</code>	The session is not returned. An <code>UnauthorizedSessionRequestException</code> error is thrown*	The session is not returned. An <code>UnauthorizedSessionRequestException</code> error is thrown*

* A `com.ibm.websphere.servlet.session.UnauthorizedSessionRequestException` error is thrown to the servlet.

Session management support

WebSphere Application Server provides facilities, grouped under the heading *Session Management*, that support the `javax.servlet.http.HttpSession` interface described in the Servlet API specification.

In accordance with the Servlet 2.3 API specification, the Session Management facility supports session scoping by Web module. Only servlets in the same Web module can access the data associated with a

particular session. Multiple requests from the same browser, each specifying a unique Web application, result in multiple sessions with a shared session ID. You can invalidate any of the sessions that share a session ID without affecting the other sessions.

You can configure a session timeout for each Web application. A Web application timeout value of 0 (the default value) means that the invalidation timeout value from the Session Management facility is used.

When an HTTP client interacts with a servlet, the state information associated with a series of client requests is represented as an HTTP session and identified by a session ID. Session Management is responsible for managing HTTP sessions, providing storage for session data, allocating session IDs, and tracking the session ID associated with each client request through the use of cookies or URL rewriting techniques. Session Management can store session-related information in several ways:

- In application server memory (the default). This information cannot be shared with other application servers.
- In a database. This storage option is known as *database persistent sessions*.

The last two options are referred to as *distributed sessions*. Distributed sessions are essential for using HTTP sessions for failover facility. When an application server receives a request associated with a session ID that it currently does not have in memory, it can obtain the required session state by accessing the external store (database or memory-to-memory). If distributed session support is not enabled, an application server cannot access session information for HTTP requests that are sent to servers other than the one where the session was originally created. Session Management implements caching optimizations to minimize the overhead of accessing the external store, especially when consecutive requests are routed to the same application server.

Storing session states in an external store also provides a degree of fault tolerance. If an application server goes offline, the state of its current sessions is still available in the external store. This availability enables other application servers to continue processing subsequent client requests associated with that session.

Saving session states to an external location does not completely guarantee their preservation in case of a server failure. For example, if a server fails while it is modifying the state of a session, some information is lost and subsequent processing using that session can be affected. However, this situation represents a very small period of time when there is a risk of losing session information.

The drawback to saving session states in an external store is that accessing the session state in an external location can use valuable system resources. Session Management can improve system performance by caching the session data at the server level. Multiple consecutive requests that are directed to the same server can find the required state data in the cache, reducing the number of times that the actual session state is accessed in external store and consequently reducing the overhead associated with external location access.

Configuring session management by level

When you configure session management at the Web container level, all applications and the respective Web modules in the Web container normally inherit that configuration, setting up a basic default configuration for the applications and Web modules below it.

However, you can set up different configurations individually for specific applications and Web modules that vary from the Web container default. These different configurations override the default for these applications and Web modules only.

Note: When you overwrite the default session management settings on the application level, all the Web modules below that application inherit this new setting unless they too are set to overwrite these settings.

1. Open the Administrative console.
 2. Select the level that this configuration applies to:
 - For the web container level:
 - a. Click **Servers > Application Servers**.
 - b. Select a server from the list of application servers.
 - c. Under Additional Properties, click **Web Container**.
 - For the enterprise application level:
 - a. Click **Applications > Applications**.
 - b. Select an applications from the list of applications.
 - For the Web module level:
 - a. Click **Applications > Enterprise Applications**.
 - b. Select an applications from the list of applications.
 - c. Under Related Items, click **Web Modules**.
 - d. Select a Web module from the list of Web modules defined for this application.
 3. Under **Additional Properties**, click **Session Management**.
 4. Make whatever changes you need to manage sessions
 5. If you are working on the Web module or application level and want these settings to override the inherited Session Management settings, under **General Properties**, select **Override**.
 6. Click **Apply** and **Save**.
-

Session tracking options

There are several options for session tracking, depending on what sort of tracking method you want to use:

- Session tracking with cookies
- Session tracking with URL rewriting
- Session tracking with Secure Sockets Layer (SSL) information

Session tracking with cookies

Tracking sessions with cookies is the default. No special programming is required to track sessions with cookies.

Session tracking with URL rewriting

An application that uses URL rewriting to track sessions must adhere to certain programming guidelines. The application developer needs to do the following:

- Program servlets to encode URLs
- Supply a servlet or Java Server Pages (JSP) file as an entry point to the application

Using URL rewriting also requires that you enable URL rewriting in the Session Management facility.

Note: In certain cases, clients cannot accept cookies. Therefore, you cannot use cookies as a session tracking mechanism. Applications can use URL rewriting as a substitute.

Program session servlets to encode URLs

Depending on whether the servlet is returning URLs to the browser or redirecting them, include either `encodeURL()` or `encodeRedirectURL()` in the servlet code. Examples demonstrating what to replace in your current servlet code follow.

Rewrite URLs to return to the browser

Suppose you currently have this statement:

```
out.println("<a href=\"/store/catalog\">catalog<a>");
```

Change the servlet to call the `encodeURL` method before sending the URL to the output stream:

```
out.println("<a href=\"\"");  
out.println(response.encodeURL ("/store/catalog"));  
out.println(">catalog</a>");
```

Rewrite URLs to redirect

Suppose you currently have the following statement:

```
response.sendRedirect ("http://myhost/store/catalog");
```

Change the servlet to call the `encodeRedirectURL` method before sending the URL to the output stream:

```
response.sendRedirect (response.encodeRedirectURL ("http://myhost/store/catalog"));
```

The `encodeURL()` and `encodeRedirectURL()` methods are part of the `HttpServletResponse` object. These calls check to see if URL rewriting is configured before encoding the URL. If it is not configured, the calls return the original URL.

If both cookies and URL rewriting are enabled and `response.encodeURL()` or `encodeRedirectURL()` is called, the URL is encoded, even if the browser making the HTTP request processed the session cookie.

You can also configure session support to enable protocol switch rewriting. When this option is enabled, the product encodes the URL with the session ID for switching between HTTP and HTTPS protocols.

Supply a servlet or JSP file as an entry point

The entry point to an application (such as the initial screen presented) may not require the use of sessions. However, if the application in general requires session support (meaning some part of it, such as a servlet, requires session support), then after a session is created, all URLs are encoded to perpetuate the session ID for the servlet (or other application component) requiring the session support.

The following example shows how you can embed Java code within a JSP file:

```
<%  
response.encodeURL ("/store/catalog");  
%>
```

Session tracking with SSL information

No special programming is required to track sessions with Secure Sockets Layer (SSL) information.

To use SSL information, turn on **Enable SSL ID tracking** in the Session Management property page. Because the SSL session ID is negotiated between the Web browser and HTTP server, this ID cannot survive an HTTP server failure. However, the failure of an application server does not affect the SSL session ID if an external HTTP Server is present between WebSphere Application Server and the browser.

SSL tracking is supported for the IBM HTTP Server and iPlanet Web servers only. You can control the lifetime of an SSL session ID by configuring options in the Web server. For example, in the IBM HTTP Server, set the configuration variable `SSLV3TIMEOUT` to provide an adequate lifetime for the SSL session ID. An interval that is too short can cause a premature termination of a session. Also, some Web browsers might have their own timers that affect the lifetime of the SSL session ID. These Web browsers may not leave the SSL session ID active long enough to serve as a useful mechanism for session tracking. Internal Http Server of WebSphere also supports SSL Tracking.

When using the SSL session ID as the session tracking mechanism in a cloned environment, use either cookies or URL rewriting to maintain session affinity. The cookie or rewritten URL contains session affinity information that enables the Web server to properly route a session back to the same server for each request.

Configuring session tracking

To configure session tracking, complete the following:

1. Go to the appropriate level of Session Management.
2. Specify which session tracking mechanism you want to pass the session ID between the browser and the servlet:
 - To track sessions with cookies, click **Enable Cookies**.
To change the cookie settings, click **Modify**.
 - To track sessions with URL rewriting, click **Enable URL Rewriting**.
If you want to enable protocol switch rewriting, click **Enable protocol switch rewriting**.
 - To track sessions with SSL information, click **Enable SSL ID tracking**.
3. Click **Apply**.
4. Click **Save**.
5. Define the session recovery characteristics.

Serializing access to session data

The Servlet API supports concurrent access to a session in a given server instance. WebSphere Application Server provides an option to prevent the concurrent access to a session in a given server instance so that concurrent modification of a session does not occur in a given server instance. This prevention is achieved by synchronizing the requests based on session. When this feature is turned on, a session is obtained for the request before invoking the servlet and requests are synchronized by locking the session for the servlet execution time. Note that synchronization is based on the memory copy of session. So this feature cannot serialize requests across servers based on session when session affinity fails.

Restriction: Use this feature only when concurrent modification of the same session data is possible and is not desirable by the application. This feature has overhead of serializing the requests based on a session.

Do the following to synchronize session access:

1. Select the level of Session Management on which you want to serialize session access.
2. Under Serialize Session access, click **Allow serial access**.
3. In the Maximum wait time box, type the amount of time, in milliseconds, a servlet waits on a session before continuing execution. The default is 120000 milliseconds or two minutes.
4. Select **Allow access on timeout** if you want the servlet to gain access to the session and continue normal execution even if the session is still locked by another servlet. If you do not select this box, the servlet execution will abort when the session request times out.
5. Click **Apply**.
6. Click **Save**.

Session Management settings

Use this page to manage HTTP session support. This support includes specifying a session tracking mechanism, setting maximum in-memory session count, controlling overflow, and configuring session timeout.

To view this administrative console page, click **Servers > Application Servers > server_name > Web Container > Session Management**.

Overwrite Session Management

Specifies whether or not these session management settings take precedence over those normally inherited from a higher level for the current application or web module.

By default, web modules inherit session management settings from the application level above it, and applications inherit session management settings from the web container level above it.

Session tracking mechanism

Specifies a mechanism for HTTP session management.

Mechanism	Function	Default
Enable SSL ID Tracking	Specifies that session tracking uses Secure Sockets Layer (SSL) information as a session ID. Enabling SSL tracking takes precedence over cookie-based session tracking and URL rewriting.	9600 seconds

There are two parameters available if you enable SSL ID tracking: SSLV3Timeout and Secure Authentication Service (SAS). SSLV3Timeout specifies the time interval after which SSL sessions are renegotiated. This is a high setting and modification does not provide any significant impact on performance. The SAS parameter establishes an SSL connection only if it goes out of the Java Virtual Machine (JVM) to another JVM. If all the beans are co-located within the same JVM, the SSL used by SAS does not hinder performance.

These are set by editing the `sas.server.properties` and `sas.client.props` files located in the `product_installation_root\properties` directory, where `product_installation_root` is the directory where WebSphere Application Server is installed.

Enable Cookies

Specifies that session tracking uses cookies to carry session IDs. If cookies are enabled, session tracking recognizes session IDs that arrive as cookies and tries to use cookies for sending session IDs. If cookies are not enabled, session tracking uses Uniform Resource Identifier (URL) rewriting instead of cookies (if URL rewriting is enabled).

Enabling cookies takes precedence over URL rewriting. Do not disable cookies in the Session Management facility of the application server that is running the administrative application because this action causes the administrative application not to function after a restart of the server. As an alternative, run the administrative application in a separate process from your applications.

Enable URL Rewriting

Click **Modify** to change these settings. Specifies that the session management facility uses rewritten URLs to carry the session IDs. If URL rewriting is enabled, the session management facility recognizes session IDs that arrive in the URL if the `encodeURL` method is called in the servlet.

Enable Protocol Switch Rewriting

Specifies that the session ID is added to a URL when the URL requires a switch from HTTP to HTTPS or from HTTPS to HTTP. If rewriting is enabled, the session ID is required to go between HTTP and HTTPS.

Maximum in-memory session count

Specifies the maximum number of sessions to maintain in memory.

The meaning differs depending on whether you are using in-memory or distributed sessions. For in-memory sessions, this value specifies the number of sessions in the base session table. Use the `Allow Overflow` property to specify whether to limit sessions to this number for the entire Session Management facility or to allow additional sessions to be stored in secondary tables. For distributed sessions, this value specifies the size of the memory cache for sessions. When the session cache has reached its maximum size and a new session is requested, the Session Management facility removes the least recently used session from the cache to make room for the new one.

Note: Do not set this value to a number less than the maximum thread pool size for your server.

Overflow

Specifies that the number of sessions in memory can exceed the value specified by the `Max In Memory Session Count` property. This option is valid only in nondistributed sessions mode.

Session timeout

Specifies how long a session can go unused before it is no longer valid. Specify either Set timeout or No timeout. Specify the value in minutes greater than or equal to two.

The value specified in a web module deployment descriptor file takes precedence over the administrative console settings. However, the value of this setting is used as a default when the session timeout is not specified in a web module deployment descriptor. Note that to preserve performance, the invalidation timer is not accurate to the second. When the Write Frequency is time based, ensure that this value is least twice as large as the write interval.

Security integration

Specifies that when security integration is enabled, the session management facility associates the identity of users with their HTTP sessions

Serialize session access

Specifies that concurrent session access in a given server is not allowed.

Maximum wait time

Specifies the maximum amount of time a servlet request waits on an HTTP session before continuing execution. This parameter is optional and expressed in seconds. The default is 120, or 2 minutes. Under normal conditions, a servlet request waiting for access to an HTTP session gets notified by the request that currently owns the given HTTP session when the request finishes.

Allow access on timeout

Specifies whether the servlet is executed normally or aborted in the event of a timeout. If this box is checked, the servlet executes normally. If this box is not checked, the servlet execution aborts and error logs are generated.

Cookie settings

Use this page to configure cookie settings for session management.

To view this administrative console page, click **Servers > Application Servers > *server_name* > Web Container > Session Management > Enable Cookies.**

Cookie name

Specifies a unique name for the session management cookie. The servlet specification requires the name JSESSIONID. However, for flexibility this value can be configured.

Secure cookies

Specifies that the session cookies include the secure field. Enabling the feature restricts the exchange of cookies to HTTPS sessions only.

Cookie domain

Specifies the domain field of a session tracking cookie. This value controls whether or not a browser sends a cookie to particular servers. For example, if you specify a particular domain, session cookies are sent to hosts in that domain. The default domain is the server.

Cookie path

Specifies that a cookie is sent to the URL designated in the path. Specify any string representing a path on the server. "/" indicates root directory. Specify a value to restrict the paths to which the cookie will be sent. By restricting paths, you prevent the cookie from going to certain URLs on the server. If you specify the root directory, the cookie is sent no matter which path on the given server is accessed.

Cookie maximum age

Specifies the amount of time that the cookie lives on the client browser. Specify that the cookie lives only as long as the current browser session, or to a maximum age. If you choose the maximum age option, specify the age in seconds. This value corresponds to the Time to Live (TTL) value described in the Cookie specification.

Default is the current browser session which is equivalent to setting the value to -1.

Configuring session tracking for Wireless Application Protocol (WAP) devices

Most Wireless Application Protocol (WAP) devices do not support cookies. The preferred way to track sessions for WAP devices is to use URL rewriting. However on most WAP devices, the maximum allowed URL length is 128 characters. With URL rewriting, a session identifier is added to the URL itself, effectively decreasing the space available for the actual URL and the number of parameters that can be sent on a request.

To reduce the length of session identifier, you can configure key (jsessionid), session ID length and clone ID. To make these configuration changes, complete the following:

1. Open the Administrative console.
2. Click **Servers > Application Servers**.
3. Select a server from the list of application servers.
4. Under Additional Properties, click **Web Container**
5. Under Additional Properties, click **Custom Properties**.
6. Add the appropriate properties from the following list:
 - HttpSessionIdLength
 - SessionRewriteIdentifier
 - HttpSessionCloneId
 - CloneSeparatorChange
 - NoAdditionalSessionInfo
 - SessionIdentifierMaxLength
7. Click **Apply** and **Save**.

Distributed sessions

WebSphere Application Server provides the following session mechanisms in a distributed environment:

- **Database Session persistence**, where sessions are stored in the database specified.

When a session contains attributes that implement `HttpSessionActivationListener`, notification occurs anytime the session is activated (that is, session is read to the memory cache) or passivated (that is, session leaves the memory cache). Passivation can occur because of a server shutdown or when the session memory cache is full and an older session is removed from the memory cache to make room for a newer session. It is not guaranteed that a session is passivated in one application server prior to being activated in another.

Session recovery support

For session recovery support, WebSphere Application Server provides distributed session support in the form of database sessions. Use session recovery support under the following conditions:

- When the user's session data must be maintained across a server restart
- When the user's session data is too valuable to lose through an unexpected server failure

All the attributes set in a session must implement `java.io.Serializable` if the session requires external storage. In general, consider making all objects held by a session serialized, even if immediate plans do not call for session recovery support. If the Web site grows, and session recovery support becomes necessary, the transition occurs transparently to the application if the sessions only hold serialized objects. If not, a switch to session recovery support requires coding changes to make the session contents serialized.

Distributed Environment settings

Use this page to specify a type for saving a session in a distributed environment.

To view this administrative console page, click **Servers > Application Servers > *server_name* > Web Container > Session Management > Distributed Environment Settings**.

Distributed Sessions

Specifies the type of distributed environment to be used for saving sessions.

None	Specifies that the session management facility discards the session data when the server shuts down.
Database	Specifies that the session management facility stores session information in the data source specified by the data source connection settings. Click Database to change these data source settings.

Configuring for database session persistence

To configure the session management facility for database session persistence, complete the following:

1. Define a JDBC provider.
2. Create a data source pointing to an existing database, using the JDBC provider that you defined. The data source should be non-JTA, for example, non-XA enabled. Note the JNDI name of the data source. Under **Data Sources > *datasource_name* > Custom Properties**, make sure the correct database is entered for the value of the **databaseName** property. If necessary, contact your database administrator to verify the correct database name.
3. Go to the appropriate level of Session Management.
4. Click **Distributed Environment Settings**
5. Select and click **Database**.
6. Specify the Data Source JNDI name from step 2.
7. Specify the database user ID and password for accessing the database.
8. Retype the password for confirmation.
9. Configure a table space and page sizes for DB2 session databases.
10. Switch to a multirow schema.
11. Click **OK**.
12. If you want to change the tuning parameters, click **Custom Tuning Parameters** and select a setting or customize one.
13. Click **Apply**.
14. Click **Save**.

Switching to a multirow schema

By default, a single session maps to a single row in the database table used to hold sessions. With this setup, there are hard limits to the amount of user-defined, application-specific data that WebSphere Application Server can access.

1. Modify the Session Management facility properties to switch from single to multirow schema.
2. Manually drop the database table or delete all the rows in the database table that the product uses to maintain HttpSession objects.

To drop the table:

- a. Determine which data source configuration Session Management is using.
- b. In the data source configuration, look up the database name.
- c. Use the database facilities to connect to the database.
- d. Drop the SESSIONS table.

Configuring tablespace and page sizes for DB2 session databases

If you are using DB2 for session persistence, you can increase the page size to optimize performance for writing large amounts of data to the database. Page sizes of 8K, 16K, and 32K are supported.

To use a page size other than the default (4K), do the following:

1. If the SESSIONS table already exists, drop it from the DB2 database.
2. Create a new DB2 buffer pool and table space, specifying the same page size (8K, 16K or 32K) for both, and assign the new buffer pool to this table space.

```
DB2 Connect to session
DB2 CREATE BUFFERPOOL sessionBP SIZE 1000 PAGESIZE 8K
DB2 Connect reset
DB2 Connect to session
DB2 CREATE TABLESPACE sessionTS PAGESIZE 8K MANAGED BY SYSTEM
    USING ('D:\DB2\NODE0000\SQL00005\sessionTS.0') BUFFERPOOL sessionBP
DB2 Connect reset
```

Refer to DB2 product documentation for details.

3. Configure the correct table space name and page size in the Session Management facility. Page size is referred to as *row size* on the Session Management page.)

When the product is restarted, the Session Management facility creates the new SESSIONS table in the specified tablespace based on the indicated page size.

Database settings

Use this page to specify the settings for database session support.

To view this administrative console page, click **Servers > Application Servers > server_name > Web Container > Session Management > Distributed Environment Settings > Database.**

Datasource JNDI Name

Specifies the datasource description

The JNDI name of the non-XA enabled data source from which Session Management obtains database connections. For example, if the JNDI name of the datasource is "jdbc/sessions", specify "jdbc/sessions." The data source represents a pool of database connections and a configuration for that pool (such as the pool size). The data source must already exist as a configured resource in the environment.

User ID

Specifies the user ID for database access

Password

Specifies the password for database access

Confirm Password

Specifies the password a second time to ensure it recorded correctly.

DB2 Row Size

Specifies the tablespace page size configured for the sessions table, if using a DB2 database. Possible values are 4, 8, 16, and 32 kilobytes (K). The default row size is 4K.

The default row size is 4K. In DB2, it can be updated to a larger value. This can help database performance in some environments. When this value is other than 4, you must specify Table Space Name to use. For 4K pages, the Table Space Name is optional.

Table Space Name

Specifies that tablespace to be used for the sessions table.

This value is required when the DB2 Page Size is other than 4K.

Multi row schema

Specifies that each instance of application data be placed in a separate row in the database, allowing larger amounts of data to be stored for each session. This action can yield better performance in certain usage scenarios. If using multirow schema is not enabled, instances of application data can be placed in the same row.

Multirow schema considerations

IBM WebSphere Application Server supports the use of a multirow schema option in which each piece of application specific data is stored in a separate row of the database. With this setup, the total amount of data you can place in a session is now bound only by the database capacities. The only practical limit that remains is the size of the session attribute object.

The multirow schema potentially has performance benefits in certain usage scenarios, such as when larger amounts of data are stored in the session but only small amounts are specifically accessed during a given servlet processing of an HTTP request. In such a scenario, avoiding unneeded Java object serialization is beneficial to performance.

Understand that switching between multirow and single row is not a trivial proposition.

In addition to allowing larger session records, using multirow schema can yield performance benefits. However, it requires a little work to switch from single-row to multirow schema, as shown in the instructions below.

Coding considerations and test environment

Consider configuring direct single-row usage to one database and multirow usage to another database while you verify which option suits your application needs. (Do this in code by switching the data source used; then monitor performance.)

Programming issue	Application scenario
Reasons to use single-row	<ul style="list-style-type: none">You can read or write all values with just one record read and write.This takes up less space in a database because you are guaranteed that each session is only one record long.
Reasons not to use single-row	2-megabyte limit of stored data per session.

Programming issue	Application scenario
Reasons to use multirow	<ul style="list-style-type: none"> • The application can store an unlimited amount of data; that is, you are limited only by the size of the database and a 2-megabyte-per-record limit. • The application can read individual fields instead of the whole record. When large amounts of data are stored in the session but only small amounts are specifically accessed during servlet processing of an HTTP request, multirow sessions can improve performance by avoiding unneeded Java object serialization.
Reasons not to use multirow	If data is small in size, you probably do not want the extra overhead of multiple row reads when you can store everything in one row.

In the case of multirow usage, design your application data objects not to have references to each other, to prevent circular references. For example, suppose you are storing two objects A and B in the session using `HttpSession.put(..)` method, and A contains a reference to B. In the multirow case, because objects are stored in different rows of the database, when objects A and B are retrieved later, the object graph between A and B is different than stored. A and B behave as independent objects.

Memory-to-memory replication

WebSphere Application Server supports session replication to another WebSphere Application Server instance. This support is referred to as *memory-to-memory session replication*. In this mode, sessions can replicate to one or more WebSphere Application Server instances to address HTTP Session single point of failure (SPOF). This is a new alternative in IBM WebSphere Application Server, Version 5 to the existing saving of HTTP Session to a database.

The WebSphere Application Server instance in which the session is currently processed is referred to as the *owner of the session*. In a clustered environment, session affinity in the WebSphere Application Server plug-in routes the requests for a given session to the same server. If the current owner server instance of the session fails, then the WebSphere Application Server plug-in routes the requests to another appropriate server in the cluster. This server either retrieves the session from a server that has the backup copy of the session or it retrieves the session from its own backup copy table. The server now becomes the owner of the session and affinity is now maintained to this server.

When a session is created or updated in a WebSphere Application Server instance, the session is transferred (or replicated) through one of the replicator entries under the replication domain that is configured with the session management facility. This session potentially gets replicated to the WebSphere Application Server instances that are also connected to the same replicator domain. The mode and partitioning determine whether WebSphere Application Server instances in the same replication domain gets the session.

There are three possible modes. You can set up a WebSphere Application Server instance to run in:

- **Server mode:** Only store backup copies of other WebSphere Application Server sessions and not to send out copies of any session created in that particular server
- **Client mode:** Only broadcast or send out copies of the sessions it owns and not to receive backup copies of sessions from other servers
- **Both mode:** Simultaneously broadcast or send out copies of the sessions it owns and act as a backup table for sessions owned by other WebSphere Application Server instances

You can select the replication mode of server, client, or both when configuring the session management facility for memory-to-memory replication. The default is both. This storage option is controlled by the mode parameter.

With respect to mode, the following are the primary examples of memory-to-memory replication configuration:

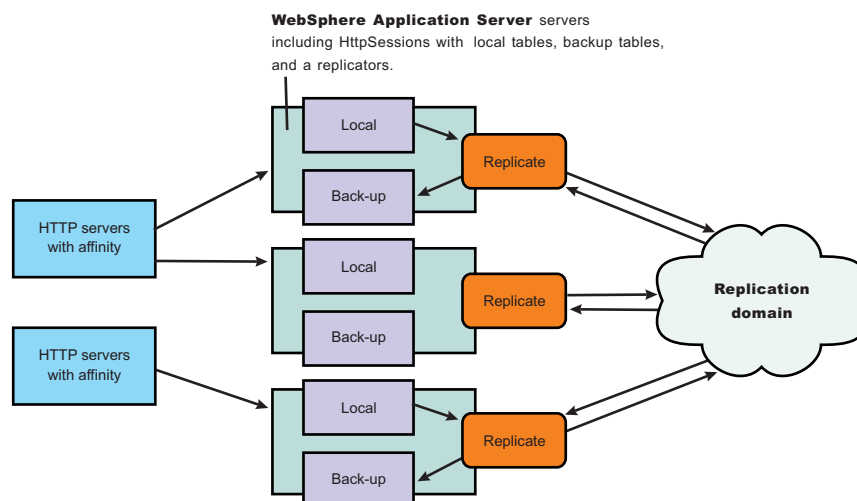
- Peer-to-peer with a local replicator
- Peer-to-peer with remote replicators
- Client/server with remote replicators
- Client/server with isolated replicators

In a cluster, by default, sessions are replicated in all the servers in the cluster that are connected to the same replicator domain. This replication can be redundant if a large number of servers exist in a cluster. The session management facility has an option to partition the servers into groups when storing sessions.

Memory-to-memory topology: Peer-to-peer function with a local replicator

The basic peer-to-peer (both client and server function, or both mode) topology is the default configuration. This configuration uses a local replicator (a replicator on the same server as the session manager) and no partitioning.

Memory to memory topology: Basic peer to peer or “both mode”



In this basic peer-to-peer topology, each server Java Virtual Machine (JVM):

- Host the Web application leveraging HTTP session
- Send out changes to the HTTP session that it owns
- Receive backup copies of the HTTP session from all of the other servers in the cluster
- House a replicator to which the session manager connects

This configuration represents the most consolidated topology, where the various system parts are collocated and requires the fewest server processes. When using this configuration, the most stable implementation is achieved when each node has equal capabilities (CPU, memory, and so on), and each handles the same amount of work.

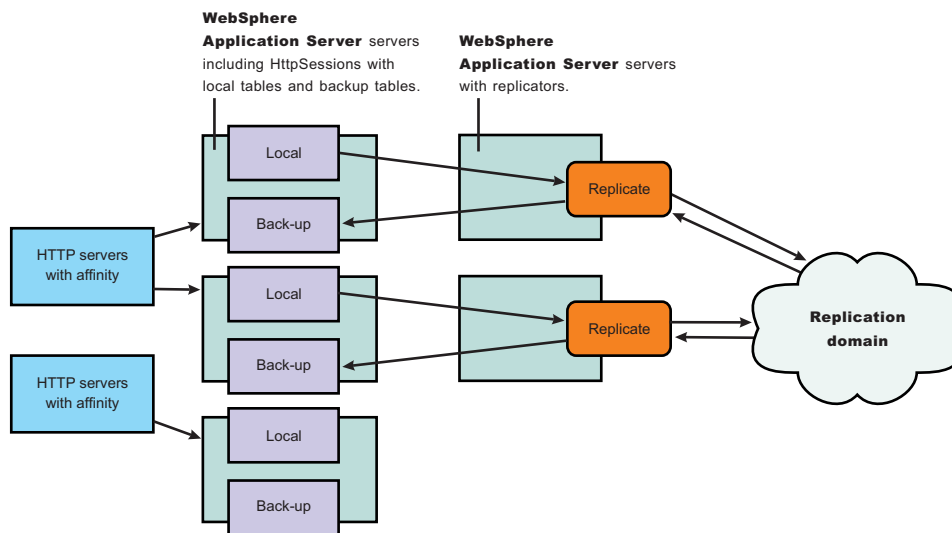
This topology is the most redundant because everyone replicates to everyone. On any failure recovery scenario, the server routed to already has a copy of the session. But as you add servers, more overhead (both CPU and memory) is needed to deal with replication. Without partitioning, scaling breaks down the

quickest with this configuration. When scaling breaks down because of the machine characteristics, the client hit rate, the network speed, and the size on average of the HTTP session. The remaining topologies: client/server with isolated replicators, client/server with remote replicators, and peer-to-peer with remote replicators help mitigate and manage the scaling characteristics of replication. Replication costs are also reduced through the use of partitioning and time based write (employed by default) in these topologies.

Memory-to-memory topology: Peer-to-peer function with remote/isolated replicators

In this topology, the replicators are separated from the clustered servers running the deployed application.

Memory to memory topology: Peer to peer with remote replicators



This method increases the number of servers used, which can be an increased cost for any installation. However, if containable, the separation of replicators produces several benefits:

- A replicator can handle the activity from session managers in multiple servers. A one-to-one correspondence between the replicators and servers can often be unnecessary; having a primary server (with a second server as a backup) is often sufficient.
- Not running the replicator in the Java Virtual Machine (JVM) running the application can free up some heap and central processing unit (CPU) for running the application.
- Putting components into separate processes means that if one of the component processes fails, this failure does not affect the other components.

Replicators can often handle multiple servers (six to seven). By using this topology, you have two configuration options:


- Use one replicator, and keep the other as a backup on standby (remember, if a replicator goes down, the session manager automatically switches to the next replicator domain).
- Split your session manager/application servers between the two replicators.

After selecting memory-to-memory replication on the session manager, you must select memory-to-memory replication to go to the detail configuration panel. You can select the replication domain and replicator for this session manager to use.

Memory-to-memory topology: Client/server function with remote replicators

The following figure depicts the client/server with remote replicators topology. There is a tier of applications servers that host web applications using HTTP Session, and these sessions are replicated out as they are created and updated. There is a second tier of servers without a web application installed, where the session manager receives updates from the replication clients. The replicators facilitating the transfer of data reside with the replication servers.

Memory to memory topology: Clients and servers with remote replicators



Benefits of the client/server with remote replicators configuration include:

- Isolation (for failure recovery): in this case we are isolating the handling of backup data from local data; aside from isolating the moving parts in case of a catastrophic failure in one of them, you again free up memory and processing in the servers processing the web application, much the same as the isolating of the replicators as showed in the topology for peer-to-peer function with remote replicators.
- Isolation (for stopping and starting): you can recycle a backup without affecting the servers running the application (when there are two or more backups, failure recovery is possible), and conversely recycle an application JVM without potentially losing that backup data for someone.
- Consolidation: there is most likely no need to have a one-to-one correspondence between servers handling backups and those processing the applications; hence, you are again reducing the number of places to which you transfer the data.
- Disparate hardware: while you run your web applications on cheaper hardware, you may have one or two more powerful computers in the back end of your enterprise that have the capacity to run a couple of session managers in replication server mode; allowing you to free up your cheaper web application hardware to process the web application.

You can define replicators on both the replication client and the server Java Virtual Machines (JVMs), however, replicators on both the clients and the servers is redundant. In the client/server topologies defining replicators only on the backup replication server JVMs is recommended.

It is good to spread out the replication clients as equally across the replicators as possible (by default they select the first replicator in the domain), because both replication servers are doing work and not acting as a hot standby.

Timing considerations: Start the replication servers first to avoid unexpected timing windows. The clients attempt to reconnect to the replication domain if you start the replication clients before the replications

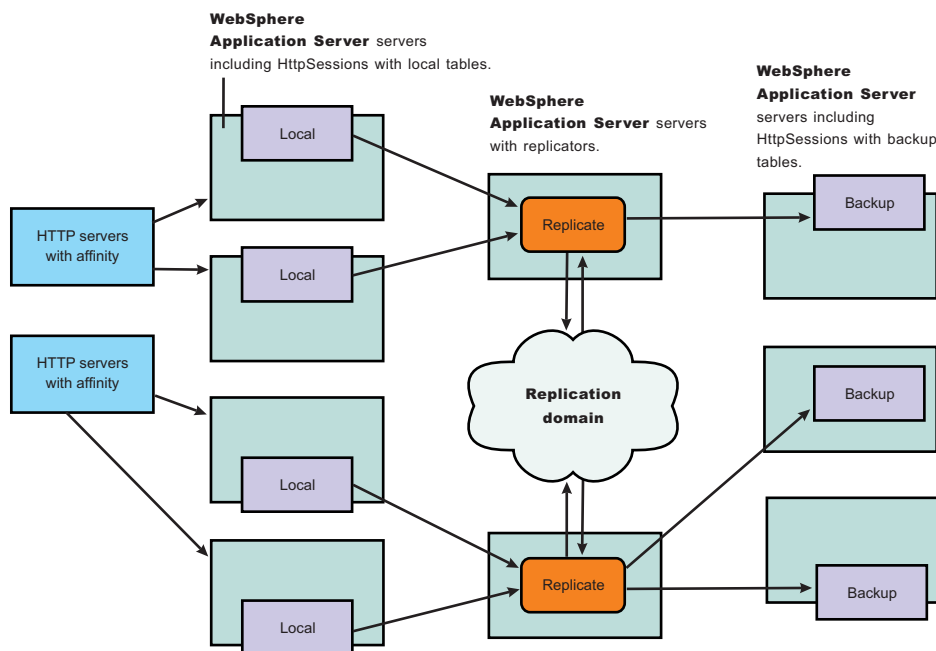
servers, even if the initial connection cannot be completed. However, if servers with the application come up, and requests on the applications occur before the replicators on the backup servers finish coming up, some expected client replication might not occur.

Memory-to-memory topology: Client/server function with isolated replicators

The following figure depicts the client/server with isolated replicators topology. In this topology, the replicators are separated from the backup servers and are on their own servers.

There is a tier of applications servers that host web applications using HTTP Session, and these sessions are replicated out as they are created and updated. There is a second tier of servers without a web application installed, where the session manager receives updates from the replication clients. The replicators facilitating the transfer of data reside with the replication servers

Memory to memory topology: Clients and servers with isolated replicators



Benefits of this topology include:

- The catastrophic failure of one server does not affect another server because of separate moving parts (the replicator and the backup server).
- You do not need a one-to-one correspondence between the replicators and the backup servers.

Although this topology is the most expensive from a Java Virtual Machine (JVM) perspective, and requires the most configuration through the administrative console, it provides the greatest degree of isolation from failure of a component, the most flexibility in starting/stopping the parts, and the most flexibility in mixing the right number of clients, servers, and replicators. Since both servers and clients need to connect to the replicators, starting the replicators removes the need for connection retries. Similarly, starting the clients after the servers, avoid timing windows where the clients send data before the servers are ready to receive them.

Memory-to-memory session partitions

In a replication domain, by default, sessions are replicated to all the servers that are connected to the same replication domain. This replication can be redundant if a large number of servers are located in a replication domain.

Group partitioning

The session management facility has an option to partition the servers into groups when storing sessions. From the session replication perspective, you can view the replicator as having n groups or partitions. The default is 10 groups (the recommended minimum for performance reasons). The number of groups is global for the entire replication domain and all the session managers connected into the replication domain through any of the replicators in that domain.

By default, the session management facility running in server mode listens to all the partitions on the replicator. When a session is created in this facility, it is written to one of the partitions on the replicator through a round-robin assignment performed by the session manager. At the session manager level, you can configure the session manager (when in either server or both mode) to listen to a subset of the group (by default, the session manager listens to all of them). Therefore, you can reduce or partition where the HTTP sessions get replicated. Also note that because a session manager listens to a subset of the groups, it does not mean that those are the only groups it replicates sessions on. The session manager round robins over all the groups being listened to by the complete set of session managers connected to the replication domain.

In a cluster environment with a large number of servers, some servers can be configured to listen on some partitions and others on a different partition. All the servers are still logically linked together through the replicator domain and you can retrieve a session in any server in the cluster. The number of groups of the replication domain is set under manage internal replication in the Environment tab. Configure which groups and partitions a specific Session Management facility listens to under session management.

Single replica partitioning

The alternative to group partitioning is to replicate the session to only one other server. When this option is chosen, a session manager dynamically picks another session manager that is connected to the same replication domain to replicate the HTTP session to during session creation. All updates to the session are only replicated to that single server. A round-robin algorithm is employed over the eligible session manager instances. The round robin will not include selection of other session managers in the same Java virtual machine (JVM), or other JVMs on the same node unless there is only 1 node connected to the replication domain. Because this is set at the replication domain level, every session manager connected to this replication domain uses single replica replication. The mode of the session manager has the same meaning, that is, clients only send out changes to one other session manager (albeit not in the same JVM or node), servers only receive, and session managers in both mode, send and receive changes.

Clustered session support

A clustered environment supports load balancing, where the workload is distributed among the application servers that compose the cluster. In a cluster environment, the same Web application must exist on each of the servers that can access the session. You can accomplish this setup by installing an application onto a cluster definition. Each of the servers in the group can then access the Web application

In a clustered environment, the Session Management facility requires an affinity mechanism so that all requests for a particular session are directed to the same application server instance in the cluster. This requirement conforms to the Servlet 2.3 specification in that multiple requests for a session cannot coexist in multiple application servers. One such solution provided by IBM WebSphere Application Server is *session affinity* in a cluster; this solution is available as part of the WebSphere Application Server plug-ins for Web servers. It also provides for better performance because the sessions are cached in memory. In clustered environments other than WebSphere Application Server clusters, you must use an affinity mechanism (for example, IBM WebSphere Edge Server affinity).

If one of the servers in the cluster fails, it is possible for the request to reroute to another server in the cluster. If distributed sessions support is enabled, the new server can access session data from the database or another WebSphere Application Server instance. You can retrieve the session data only if a new server has access to an external location from which it can retrieve the session.

Configuring memory-to-memory replication for the peer-to-peer function with a local replicator (default memory-to-memory replication)

To configure the session management facility for memory-to-memory session replication for peer-to-peer functions (both client and server function, or both mode) with a local replicator and no partitioning, complete the following steps:

1. Create an application cluster. This cluster is used to deploy the application.
 - a. Go to the Server Cluster page. Click **Servers> Clusters**.
 - b. Click **New**.
 - c. Type a cluster name for this application cluster.
 - d. Define a replication domain. Select the **Create Replication Domain for this cluster** check box.
 - e. Click **Next**.
 - f. Define each cluster member server. Type a cluster member name.
 - g. Select the **Create Replication Entry in this Server** check box. This step adds a replicator into this new server.
 - h. Click **Apply**. Repeat steps **f** through **h** for each server created in this cluster.
 - i. Click **Next** and review the summary of changes.
 - j. Click **Finish** to complete the configuration.

You have now created a cluster that contains the deployed application and the replication domain.

2. Enable memory-to-memory session replication for each server.
 - a. Go to the appropriate level of session management for the Web container level. Click **Application Servers> server_name> Web Container> Session Management**
 - b. Click **Distributed Environment Settings** under Additional Properties.
 - c. If a replicator has been created and is associated with this or another server, select **Memory to Memory Replication**. If no replicator is already created for any server, a note appears stating this. In this case, click **Memory to Memory Replication**. The Internal Replication Domain panels appear. These panels lead you through the process of creating a replication domain and a replicator. After you create a replication domain and a replicator, begin again with step **2a**.
 - d. Click **Apply**.
 - e. Repeat these steps for each server. Click **OK**.
3. If you want to change the tuning parameters, click **Custom Tuning Parameters** and select a setting or customize one.

Note: Using the default tuning parameter custom settings, which specifies time based write interval of 10 seconds, may result in data loss when an application server in your cluster fails. However, this is just a small opportunity for lost data when compared to the significant improvement in performance.

- a. Click **Apply**.
- b. Click **Save**.

Memory-to-memory sessions settings

Use this page to configure memory-to-memory sessions.

To view the Memory-to-memory Sessions page, click **Servers > Application Servers > *server_name* > Web Container > Session Management > Distributed Environment Settings > Memory to Memory Replication**.

Replication

Use one of these options to select the replicator for the session management facility to use for memory-to-memory replication.

Select replicator from the following domain

Specifies a replicator from a replication domain.

Click this option, and select a domain. Then click **Apply** to show all the replicators in that domain. Select the replicator in that domain.

Select replicator from another domain

Specifies a replicator in another domain that is not managed by this cell.

Click this option, and then type the IP address and the port number of the replicator.

Listen to partition groups

Specifies partitions on which the session management facility listens when running in Server mode for memory-to-memory session replication. The number of partitions is defined in the Manage Internal Replication pages. For Server mode and Both mode, the default is all partitions.

Runtime mode

Select the mode in which this server has to run: Both, Client and Server. The mode implies whether data is only sent (client), only received (server), or both. The default is both.

Configuring memory-to-memory replication for peer-to-peer functions with remote/isolated replicators

To configure the session management facility for memory-to-memory replication for peer-to-peer functions with remote replicators, complete the following steps:

1. Create an application cluster. This cluster is used to deploy the application.
 - a. Go to the Server Cluster page. Click **Servers> Clusters**.
 - b. Click **New**.
 - c. Type a cluster name for the application cluster.
 - d. Click **Next**.
 - e. Define each cluster member server. Type a cluster member name.
 - f. Click **Apply**. Complete steps **e** and **f** for each server created in this cluster.
 - g. Click **Next** and review the summary of changes.
 - h. Click **Finish** to complete the configuration.

You have now created a cluster that contains the deployed application.

2. Create a replicator cluster. The replication domain is defined in this replicator cluster.
 - a. Go to the Server Cluster page. Click **Servers> Clusters**.
 - b. Click **New**.
 - c. Type a cluster name for this replicator cluster.
 - d. Define a replication domain. Select the **Create Replication Domain for this cluster** check box.
 - e. Click **Next**.

- f. Associate replicators for each server. Select the **Create Replication Entry in this Server** check box.
- g. Click **Apply**. Complete steps **f** and **g** for each server created in this cluster.
- h. Click **Next** and review the summary of changes.
- i. Click **Finish** to complete the configuration.

You have now created a cluster that contains the replication domain. Do not deploy the application on this cluster.

3. Enable memory-to-memory session replication for each server.
 - a. Go to the appropriate level of Session Management for the Web container level. Click **Application Servers> Server_Name> Web Container> Session Management**
 - b. Click **Distributed Environment Settings** under Additional Properties.
 - c. If a replicator is already created and associated with this or another server, select **Memory to Memory Replication**. If no replicator has been created for any server, a note appears. In this case, click **Memory to Memory Replication**. The Internal Replication Domain panels appear. These panels lead you through the process of creating a replication domain and a replicator. After you create a replication domain and a replicator, begin again with step **3a**.
 - d. Click **Apply**.
 - e. Repeat these steps for each server. Click **OK**.
4. If you want to change the tuning parameters, click **Custom Tuning Parameters** and select a setting or customize one.

Note: Using the default tuning parameter custom settings, which specifies time based write interval of 10 seconds, may result in data loss when an application server in your cluster fails. However, this is just a small opportunity for lost data when compared to the significant improvement in performance.

- a. Click **Apply**.
- b. Click **Save**.

Configuring memory-to-memory replication for the client/server function using isolated replicators

To configure the session management facility for memory-to-memory session replication with replicators separated from the backup servers, complete the following steps:

1. Create an application cluster. This cluster is used to deploy the application.
 - a. Go to the Server Cluster page. Click **Servers> Clusters**.
 - b. Click **New**.
 - c. Type a cluster name for this application cluster.
 - d. Click **Next**.
 - e. Define each cluster member server. Type a cluster member name.
 - f. Click **Apply**. Complete steps **e** and **f** for each server created in this cluster.
 - g. Click **Next** and review the summary of changes.
 - h. Click **Finish** to complete the configuration.

Do not create a replication domain or replicators on the member server for the application cluster.

2. Create a cluster of session manager replication servers (backup cluster).
 - a. Go to the Server Cluster page. Click **Servers> Clusters**.
 - b. Click **New**.
 - c. Type a cluster name for the session manager of replication servers.

- d. Click **Next**.
- e. Define each cluster member server. Type a cluster member name.
- f. Define a replicator on each cluster member server. Select the **Create replication entry in this server** check box.
- g. Click **Apply**. Complete steps **e** through **g** for each server created in this cluster.
- h. Click **Next** and review the summary of changes.
- i. Click **Finish** to complete the configuration.

This step creates a cluster of session manager replication servers and creates a replicator on each cluster member server.

3. Create a replication cluster. The replication cluster defines the replication domain and each of the replicators.
 - a. Go to the Server Cluster page. Click **Servers> Clusters**.
 - b. Click **New**.
 - c. Type a replication cluster name.
 - d. Define a replication domain. Select the **Create Replication Domain for this cluster** check box.
 - e. Click **Next**.
 - f. Define each replicator. Type a cluster member name.
 - g. Select the **Create Replication Entry in this Server** check box.
 - h. Click **Apply**. Complete steps **f** through **h** for each server created in this cluster.
 - i. Click **Next** and review the summary of changes.
 - j. Click **Finish** to complete the configuration.
4. Enable memory-to-memory session replication for each cluster member server in the application cluster and the backup cluster. Memory-to-memory session replication does not need enablement in the replication cluster; these servers act as replicators only.
 - a. Go to the appropriate level of session management for the web container level. Click **Application Server > application_cluster_member_server** or **backup_cluster_member_server > Web Container > Session Management**.
 - b. Click **Distributed Environment Settings** under Additional Properties.
 - c. Click **Memory to Memory Replication** . Verify that each session manager (both the replication clients and the servers) point to a replicator from the replication cluster, distributing them as evenly as possible.
 - d. Click **Apply**.
 - e. Repeat these steps for each server. Click **OK**.

This step points the servers to the replicator of your choice.

5. If you want to change the tuning parameters, click **Custom Tuning Parameters** and select a setting or customize one.

Note: Using the default tuning parameter custom settings, which specifies time based write interval of 10 seconds, may result in data loss when an application server in your cluster fails. However, this is just a small opportunity for lost data when compared to the significant improvement in performance.

- a. Click **Apply**.
- b. Click **Save**.

Configuring memory-to-memory replication for the client/server function using remote replicators

To configure the session management facility for memory-to-memory replication for clients and servers with remote replicators, complete the following steps:

1. Create an application cluster. This cluster is used to deploy the application.
 - a. Go to the Server Cluster page. Click **Servers> Clusters**.
 - b. Click **New**.
 - c. Type a cluster name for this application cluster.
 - d. Click **Next**.
 - e. Define each cluster member server. Type a cluster member name.
 - f. Click **Apply**. Complete steps **e** and **f** for each server created in this cluster.
 - g. Click **Next** and review the summary of changes.
 - h. Click **Finish** to complete the configuration.

Do not create a replication domain or replicators on the member server for the application cluster. You have now created a cluster that contains the replication domain.

2. Create a cluster of session manager replication servers (backup cluster).
 - a. Go to the Server Cluster page. Click **Servers> Clusters**.
 - b. Click **New**.
 - c. Type a cluster name for the cluster of session manager replication servers.
 - d. Define a replication domain. Select the **Create Replication Domain for this cluster** check box.
 - e. Click **Next**.
 - f. Define each cluster member server. Type a cluster member name.
 - g. Define a replicator on each cluster member server. Select the **Create Replication Entry in this Server** check box.
 - h. Click **Apply**. Complete steps **f** through **h** for each server created in this cluster.
 - i. Click **Next** and review the summary of changes.
 - j. Click **Finish** to complete the configuration.

This step creates a cluster of session manager replication servers and associates a replication domain with that cluster. This step also creates a replicator on each cluster member server.

3. Enable memory-to-memory session replication for each cluster member server in the application cluster.
 - a. Go to the appropriate level of session management for the Web container level. Click **Application Server > application_cluster_member_server > Web Container > Session Management**.
 - b. Click **Distributed Environment Settings** under Additional Properties.
 - c. If a replicator is already created and is associated with this or another server, select **Memory to Memory Replication**. If no replicator has been created for any server, a note appears. In this case, click **Memory to Memory Replication**. The Internal Replication Domain panels appear. These panels lead you through the process of creating a replication domain and a replicator. After you create a replication domain and a replicator, begin again with step **2a**.
 - d. Click **Apply**.
 - e. Repeat these steps for each server. Click **OK**.

This step points the servers to the replicator of your choice.

4. If you want to change the tuning parameters, click **Custom Tuning Parameters** and select a setting or customize one.

Note: Using the default tuning parameter custom settings, which specifies time based write interval of 10 seconds, may result in data loss when an application server in your cluster fails. However, this is just a small opportunity for lost data when compared to the significant improvement in performance.

- a. Select the replication domain
- b. Select the replicator to which the member server connects.
- c. Ensure that the member server node is client-only. Verify that the remainder of the member servers are in client-only mode.
- d. Click **Apply**.
- e. Click **Save**.

Tuning session management

IBM WebSphere Application Server session support has features for tuning session performance and operating characteristics, particularly when sessions are configured in a distributed environment. These options support the administrator flexibility in determining the performance and failover characteristics for their environment.

The table summarizes the features, including whether they apply to sessions tracked in memory, in a database, with memory-to-memory replication, or all. Click a feature for details about the feature. Some features are easily manipulated using administrative settings; others require code or database changes.

Feature or option	Goal	Applies to sessions in memory, database, or memory-to-memory
Write frequency	Minimize database write operations.	Database and Memory-to-Memory
Session affinity	Access the session in the same application server instance.	All
Multirow schema	Fully utilize database capacities.	Database
Base in-memory session pool size	Fully utilize system capacity without overburdening system.	All
Write contents	Allow flexibility in determining what session data to write	Database and Memory-to-Memory
Scheduled invalidation	Minimize contention between session requests and invalidation of sessions by the Session Management facility. Minimize write operations to database for updates to last access time only.	Database and Memory-to-Memory
Tablespace and row size	Increase efficiency of write operations to database.	Database (DB2 only)

Configuring scheduled invalidation

Instead of relying on the periodic invalidation timer that runs on an interval based on the session timeout parameter, you can set specific times for the session management facility to scan for invalidated sessions in a distributed environment. When used with distributed sessions, this feature has the following benefits:

- You can schedule the scan for invalidated sessions for times of low application server activity, avoiding contention between invalidation scans of database or another WebSphere Application Server instance and read and write operations to service HTTP session requests.
- Significantly fewer external write operations can occur when running with the End of Service Method Write mode because the last access time of the session does not need to be written out on each HTTP request. (Manual Update options and Time Based Write options already minimize the writing of the last access time.)

Usage considerations

- With scheduled invalidation configured, HttpSession timeouts are not strictly enforced. Instead, all invalidation processing is handled at the configured invalidation times.
- HttpSessionBindingListener processing is handled at the configured invalidation times unless the HttpSession.invalidate() method is explicitly called.
- The HttpSession.invalidate() method immediately invalidates the session from both the session cache and the external store.
- The periodic invalidation thread still runs with scheduled invalidation. If the current hour of the day does not match one of the configured hours, sessions that have exceeded the invalidation interval are removed from cache, but not from the external store. Another request for that session results in returning that session back into the cache.
- When the periodic invalidation thread runs during one of the configured hours, all sessions that have exceeded the invalidation interval are invalidated by removal from both the cache and the external store.
- The periodic invalidation thread can run more than once during an hour and does not necessarily run exactly at the top of the hour.
- If you specify the interval for the periodic invalidation thread using the HttpSessionReaperPollInterval custom property, do not specify a value of more than 3600 seconds (1 hour) to ensure that invalidation processing happens at least once during each hour.

Configuring write contents

In Session Management, you can configure which session data is written to the database or to another WebSphere instance, depending on whether you are using database persistent sessions or memory to memory replication. This flexibility allows for fewer code changes for the JSP writer when the application will be operating in a clustered environment. The following options are available in Session Management for tuning what is to be written back:

- Write changed (the default) - Write only session data properties that have been updated through setAttribute() and removeAttribute() method calls.
- Write all - Write all session data properties.

The **Write all** setting might benefit servlet and JSP writers who change Java objects' states that reside as attributes in HttpSession and do not call HttpSession.setAttribute().

However, the use of **Write all** could result in more data being written back than is necessary. If this situation applies to you, consider combining the use of **Write all** with **Time-based write** to boost performance overall. As always, be sure to evaluate the advantages and disadvantages for your installation.

With either Write Contents setting, when a session is first created, complete session information is written, including all of the objects bound to the session. When using database session persistence, in subsequent session requests, what is written to the database depends on whether a single-row or multirow schema has been set for the session database, as follows:

Write Contents setting	Behavior with single-row schema	Behavior with multirow schema
Write changed	If any session attribute is updated, all objects bound to the session are written.	Only the session data modified through setAttribute() or removeAttribute() calls is written.
Write all	All bound session attributes are written.	All session attributes that currently reside in the cache are written. If the session has never left the cache, all session attributes are written.

1. Go to the appropriate level of Session Management.
2. Click Distributed Environment Settings

3. Click Custom Tuning Parameters.
4. Select Custom Settings, and click Modify.
5. Select the appropriate write contents setting.

Configuring write frequency

In the Session Management facility, you can configure the frequency for writing session data to the database or to a WebSphere instance, depending on whether you use database distributed sessions or memory-to-memory replication. This flexibility enables you to weigh session performance gains against varying degrees of failover support. The following options are available in the Session Management facility for tuning write frequency:

- **End of service servlet**- Write session data at the end of the servlet `service()` method call.
- **Manual update**- Write session data only when the servlet calls the `IBMSession.sync()` method.
- **Time based** (the default) - Write session data at periodic intervals, in seconds (called the *write interval*).

When a session is first created, session information is always written at the end of the `service()` call.

Using the time based write or manual update options can result in loss of data in failover scenarios since the backup copy of the session in the persistent store (for example, a database or another JVM) may not be in sync with the session in the session cache.

Base in-memory session pool size

The base in-memory session pool size number has different meanings, depending on session support configuration:

- With in-memory sessions, session access is optimized for up to this number of sessions.
- With distributed sessions (meaning, when sessions are stored in a database or in another WebSphere Application Server instance); it also specifies the cache size and the number of last access time updates saved in manual update mode.

For distributed sessions, when the session cache has reached its maximum size and a new session is requested, the Session Management facility removes the least recently used session from the cache to make room for the new one.

General memory requirements for the hardware system, and the usage characteristics of the e-business site, determines the optimum value.

Note that increasing the base in-memory session pool size can necessitate increasing the heap sizes of the Java processes for the corresponding WebSphere Application Servers.

Overflow in nondistributed sessions

By default, the number of sessions maintained in memory is specified by base in-memory session pool size. If you do not wish to place a limit on the number of sessions maintained in memory and allow overflow, set `overflow` to *true*.

Allowing an unlimited amount of sessions can potentially exhaust system memory and even allow for system sabotage. Someone could write a malicious program that continually hits your site and creates sessions, but ignores any cookies or encoded URLs and never utilizes the same session from one HTTP request to the next.

When overflow is disallowed, the Session Management facility still returns a session with the `HttpServletRequest.getSession(true)` method when the memory limit is reached, and this is an invalid session that is not saved.

With the WebSphere Application Server extension to HttpSession, com.ibm.websphere.servlet.session.IBMSession, an isOverflow() method returns *true* if the session is such an invalid session. An application can check this status and react accordingly.

Controlling write operations

You can manually control when modified session data is written out to the database or to another WebSphere Application Server instance by using the sync() method in the com.ibm.websphere.servlet.session.IBMSession interface, which extends the javax.servlet.http.HttpSession interface. By calling the sync() method from the service() method of a servlet, you send any changes in the session to the external location. When *manual update* is selected as the write frequency mode, session data changes are written to an external location only if the application calls the sync() method. If the sync() method is not called, session data changes are lost when a session object leaves the server cache. When *end of service servlet* or *time based* is the write frequency mode, the session data changes are written out whenever the sync() method is called. If the sync() method is not called, changes are written out at the end of service method or on a time interval basis based on the write frequency mode selected.

```
IBMSession iSession = (IBMSession) request.getSession();
iSession.setAttribute("name", "Bob");

//force write to external store
iSession.sync( )
```

Tuning parameter settings

Use this page to set tuning parameters for distributed sessions.

To view this administrative console page, click **Servers > Application Servers > server_name > Web Container > Session Management > Distributed Environment Settings > Custom Tuning Parameters**.

Tuning Level

Specifies that the session management facility provides certain predefined settings that affect performance.

Select one of these predefined settings or customize a setting. To customize a setting, select one of the predefined settings that comes closest to the setting desired, click **Custom settings**, make your changes, and then click **OK**.

Very high (optimize for performance)

Write frequency	Time based
Write interval	300 seconds
Write contents	Only updated attributes
Schedule sessions cleanup	true
First time of day default	0
Second time of day default	2

High

Write frequency	Time based
Write interval	300 seconds
Write Contents	All session attributes
Schedule sessions cleanup	false

Medium

Write frequency	End of servlet service
Write Contents	Only updated attributes
Schedule sessions cleanup	false

Low (optimize for failover)

Write frequency	End of servlet service
Write Contents	All session attributes
Schedule sessions cleanup	false

Custom settings

Write frequency default	Time based
Write interval default	10 seconds
Write contents default	All session attributes
Schedule sessions cleanup default	false

Tuning parameter custom settings

Use this page to customize tuning parameters for distributed sessions.

To view this administrative console page, click **Servers > Application Servers > *server_name* > Web Container > Session Management > Distributed Environment Settings > Custom Tuning Parameters > Custom settings**.

Write frequency

Specifies when the session is written to the persistent store.

End of servlet service	A session writes to a database or another WebSphere Application Server instance after the servlet completes execution.
Manual update	A programmatic sync on the IBMSession object is required to write the session data to the database or another WebSphere Application Server instance.
Time based	Session data writes to the database or another WebSphere Application Server instance based on the specified Write Interval value. Default: 10 seconds

Write contents

Specifies whether updated attributes are only written to the external location or all of the session attributes are written to the external location, regardless of whether or not they changed. The external location can be either a database or another application server instance.

Only updated attributes	Only updated attributes are written to the persistent store.
All session attribute	All attributes are written to the persistent store.

Schedule sessions cleanup

Specifies when to clean the invalid sessions from a database or another application server instance.

Specify distributed sessions cleanup schedule

Enables the scheduled invalidation process for cleaning up the invalidated HTTP sessions from the external location. Enable this option to reduce the number of updates to a database or another application server instance required to keep the HTTP sessions alive. When this option is not enabled, the invalidator process runs every few minutes to remove invalidated HTTP sessions.

When this option is enabled, specify the two hours of a day for the process to clean up the invalidated sessions in the external location. Specify the times when there is the least activity in the application servers. An external location can be either a database or another application server instance.

First Time of Day (0 - 23)

Indicates the first hour, in Greenwich Mean Time (GMT), during which the invalidated sessions are cleared from the external location. Specify this value as a positive integer between 0 and 23. This value is valid only when schedule invalidation is enabled.

Second Time of Day (0 - 23)

Indicates the second hour, in Greenwich Mean Time (GMT), during which the invalidated sessions are cleared from the external location. Specify this value as a positive integer between 0 and 23. This value is valid only when schedule invalidation is enabled.

Best practices for using HTTP Sessions

- **Enable Security integration for securing HTTP sessions**

HTTP sessions are identified by session IDs. A session ID is a pseudo-random number generated at the runtime. Session hijacking is a known attack HTTP sessions and can be prevented if all the requests going over the network are enforced to be over a secure connection (meaning, HTTPS). But not every configuration in a customer environment enforces this constraint because of the performance impact of SSL connections. Due to this relaxed mode, HTTP session is vulnerable to hijacking and because of this vulnerability, WebSphere Application Server has the option to tightly integrate HTTP sessions and WebSphere Application Server security. Enable security in WebSphere Application Server so that the sessions are protected in a manner that only users who created the sessions are allowed to access them.

- **Release HttpSession objects using `javax.servlet.http.HttpSession.invalidate()` when finished.**

HttpSession objects live inside the Web container until:

- The application explicitly and programmatically releases it using the `javax.servlet.http.HttpSession.invalidate()` method; quite often, programmatic invalidation is part of an application logout function.
- WebSphere Application Server destroys the allocated HttpSession when it expires (default = 1800 seconds or 30 minutes). The WebSphere Application Server can only maintain a certain number of HTTP sessions in memory based on Session Management settings. In case of distributed sessions, when maximum cache limit is reached in memory, the Session Management facility removes the least recently used (LRU) one from cache to make room for a session.

- **Avoid trying to save and reuse the HttpSession object outside of each servlet or JSP file.**

The HttpSession object is a function of the HttpServletRequest (you can get it only through the `req.getSession()` method), and a copy of it is valid only for the life of the `service()` method of the servlet or JSP file. You *cannot* cache the HttpSession object and refer to it outside the scope of a servlet or JSP file.

- **Implement the `java.io.Serializable` interface when developing new objects to be stored in the HTTP session.**

This action allows the object to properly serialize when using distributed sessions. Without this extension, the object cannot serialize correctly and throws an error. An example of this follows:

```
public class MyObject implements java.io.Serializable {...}
```

Make sure all instance variable objects that are not marked transient are serializable.

- **The HttpSession API does not dictate transactional behavior for sessions.**

Distributed HttpSession support does not guarantee transactional integrity of an attribute in a failover scenario or when session affinity is broken. Use transactionally aware resources like enterprise Java beans to guarantee the transaction integrity required by your application.

- **Ensure the Java objects you add to a session are in the correct class path.**

If you add Java objects to a session, place the class files for those objects in the correct classpath (the Application Classpath if utilizing sharing across Web modules in an enterprise application, or the WebModule Classpath if using the Servlet 2.2-complaint session sharing) or in the directory containing other servlets used in WebSphere Application Server. In the case of session clustering, this action applies to every node in the cluster.

Because the HttpSession object is shared among servlets that the user might access, consider adopting a site-wide naming convention to avoid conflicts.

- **Avoid storing large object graphs in the HttpSession object.**

In most applications each servlet only requires a fraction of the total session data. However, by storing the data in the HttpSession object as one large object, an application forces WebSphere Application Server to process all of it each time.

- **Utilize Session Affinity to help achieve higher cache hits in the WebSphere Application Server.**

WebSphere Application Server has functionality in the HTTP Server plug-in to help with session affinity. The plug-in will read the cookie data (or encoded URL) from the browser and helps direct the request to the appropriate application or clone based on the assigned session key. This functionality increases use of the in-memory cache and reduces hits to the database or another WebSphere Application Server instance

- **Maximize use of session affinity and avoid breaking affinity.**

Using session affinity properly can enhance the performance of the WebSphere Application Server. Session affinity in the WebSphere Application Server environment is a way to maximize the in-memory cache of session objects and reduce the amount of reads to the database or another WebSphere Application Server instance. Session affinity works by caching the session objects in the server instance of the application with which a user is interacting. If the application is deployed in multiple servers of a server group, the application can direct the user to any one of the servers. If the users starts on server1 and then comes in on server2 a little later, the server must write all of the session information to the external location so that the server instance in which server2 is running can read the database. You can avoid this database read using session affinity. With session affinity, the user starts on server1 for the first request; then for every successive request, the user is directed back to server1. Server1 has to look only at the cache to get the session information; server1 never has to make a call to the session database to get the information.

You can improve performance by not breaking session affinity. Some suggestions to help avoid breaking session affinity are:

- Combine all Web applications into a single application server instance, if possible, and use modeling or cloning to provide failover support.
- Create the session for the frame page, but do not create sessions for the pages within the frame when using multiframed JSP files. (See discussion later in this topic.)

- **When using multi-framed pages, follow these guidelines:**

- Create a session in only one frame or before accessing any frame sets. For example, assuming there is no session already associated with the browser and a user accesses a multi-framed JSP file, the browser issues concurrent requests for the JSP files. Because the requests are not part of any session, the JSP files end up creating multiple sessions and all of the cookies are sent back to the browser. The browser honors only the last cookie that arrives. Therefore, only the client can retrieve the session associated with the last cookie. Creating a session before accessing multi-framed pages that utilize JSP files is recommended.

- By default, JSPs get a HttpSession using request.getSession(true) method. So by default JSPs create a new session if none exists for the client. Each JSP page in the browser is requesting a new session, but only one session is used per browser instance. A developer can use `<% @ page session="false" %>` to turn off the automatic session creation from the JSP files that will not access the session. Then if the page needs access to the session information, the developer can use `<% HttpSession session = javax.servlet.http.HttpServletRequest.getSession(false); %>` to get the already existing session that was created by the original session creating JSP file. This action helps prevent breaking session affinity on the initial loading of the frame pages.
- Update session data using only one frame. When using framesets, requests come into the HTTP server concurrently. Modifying session data within only one frame so that session changes are not overwritten by session changes in concurrent frameset is recommended.
- Avoid using multi-framed JSP files where the frames point to different Web applications. This action results in losing the session created by another Web application because the JSESSIONID cookie from the first Web application gets overwritten by the JSESSIONID created by the second Web application.
- **Secure all of the pages (not just some) when applying security to servlets or JSP files that use sessions with security integration enabled, .**

When it comes to security and sessions, it is all or nothing. It does not make sense to protect access to session state only part of the time. When security integration is enabled in the Session Management facility, all resources from which a session is created or accessed must be either secured or unsecured. You cannot mix secured and unsecured resources.

The problem with securing only a couple of pages is that sessions created in secured pages are created under the identity of the authenticated user. Only the same user can access sessions in other secured pages. To protect these sessions from use by unauthorized users, you cannot access these sessions from an unsecure page. When a request from an unsecure page occurs, access is denied and an `UnauthorizedSessionRequestException` error is thrown. (`UnauthorizedSessionRequestException` is a runtime exception; it is logged for you.)

- **Use manual update and either the sync() method or time-based write in applications that read session data, and update infrequently.**

With `END_OF_SERVICE` as write frequency, when an application uses sessions and anytime data is read from or written to that session, the `LastAccess` time field updates. If database sessions are used, a new write to the database is produced. This activity is a performance hit that you can avoid using the Manual Update option and having the record written back to the database only when data values update, not on every read or write of the record.

To use manual update, turn it on in the Session Management Service. (See the tables above for location information.) Additionally, the application code must use the `com.ibm.websphere.servlet.session.IBMSession` class instead of the generic `HttpSession`. Within the `IBMSession` object there is a method called `sync()`. This method tells the WebSphere Application Server to write the data in the session object to the database. This activity helps the developer to improve overall performance by having the session information persist only when necessary.

Note: An alternative to using the manual updates is to utilize the timed updates to persist data at different time intervals. This action provides similar results as the manual update scheme.

- Implement the following suggestions to achieve high performance:
 - If your applications do not change the session data frequently, use Manual Update and the `sync()` function (or timed interval update) to efficiently persist session information.
 - Keep the amount of data stored in the session as small as possible. With the ease of using sessions to hold data, sometimes too much data is stored in the session objects. Determine a proper balance of data storage and performance to effectively use sessions.
 - If using database sessions, use a dedicated database for the session database. Avoid using the application database. This helps to avoid contention for JDBC connections and allows for better database performance.
 - If using memory to memory sessions, define replicators only on the servers and have the client attach to server replicator.

- If using memory to memory sessions, employ partitioning (either group or single replica) as your clusters grow in size and scaling decreases.
 - Verify that you have the latest fix packs for the WebSphere Application Server.
 - Utilize the following tools to help monitor session performance.
 - Run the `com.ibm.servlet.personalization.sessiontracking.IBMTrackerDebug` servlet. - To run this servlet, you must have the servlet invoker running in the Web application you want to run this from. Or, you can explicitly configure this servlet in the application you want to run.
 - Use the WebSphere Application Server Resource Analyzer which comes with WebSphere Application Server to monitor active sessions and statistics for the WebSphere Application Server environment.
 - Use database tracking tools such as "Monitoring" in DB2. (See the respective documentation for the database system used.)
-

Managing HTTP sessions: Resources for learning:

Use the following links to find relevant supplemental information about HTTP sessions. The information resides on IBM and non-IBM Internet sites, whose sponsors control the technical accuracy of the information.

These links are provided for convenience. Often, the information is not specific to the IBM WebSphere Application Server product, but is useful all or in part for understanding the product. When possible, links are provided to technical papers and Redbooks that supplement the broad coverage of the release documentation with in-depth examinations of particular product areas.

View links to additional information about:

Programming model and decisions

- Best practices
- HTTP Session Persistence Best Practices
- Improving session persistence performance with DB2
- Persistent client state HTTP cookies specification

Programming instructions and examples

- Java Servlet documentation, tutorials, and examples site

Programming specifications

- Java Servlet 2.3 API specification download site
- J2EE 1.3 specification download site

Chapter 4. Using enterprise beans in applications

1. Design a J2EE application and the enterprise beans that it needs. See "Resources for learning" for links to design information that is specific to enterprise beans.
2. Develop any enterprise beans that your application will use.
3. Prepare for assembly. For your EJB 2.x-compliant entity beans, decide on an appropriate access intent policy.
4. Assemble the beans using the Assembly Toolkit into one or more EJB modules. This includes setting security. (See "Securing enterprise bean applications using the Assembly Toolkit" in the Information Center.)
5. Assemble the modules into a J2EE application using the Assembly Toolkit.
6. **5.1+** For a given application server, update the EJB container configuration if needed for the application to be deployed, and determine if you want to batch commands batch commands or defer commands for container managed persistence.
7. Deploy the application in an application server.
8. Test the modules.
 - As needed, debug problems with the container. See "Enterprise bean and EJB container troubleshooting tips" in the Information Center.
 - Debug access and deployment problems. See "Cannot access an enterprise bean from a servlet, a JSP file, a stand-alone program, or another client" and "Errors deploying enterprise beans" in the Information Center.
9. Assemble the production application using the Assembly Toolkit.
10. Deploy the application to a production environment.
11. Manage the application:
 - a. Manage installed EJB modules. After an application has been installed, you can manage its EJB modules individually through administrative console settings.
 - b. Manage other aspects of the J2EE application.
12. Update the module and redeploy it using the Assembly Toolkit.
13. Tune the performance of the application. See Best practices for developing enterprise beans.

Enterprise beans

An enterprise bean is a Java component that can be combined with other resources to create J2EE applications. There are three types of enterprise beans, *entity* beans, *session* beans, and *message-driven* beans.

All beans reside in EJB containers, which provide an interface between the beans and the application server on which they reside.

Entity beans store permanent data. Entity beans with container-managed persistence (CMP) require connections to a form of persistent storage. This storage might be a database, an existing legacy application, a file, or another type of persistent storage. Entity beans with bean-managed persistence manage permanent data in whichever manner is defined in the bean code. For example, they can write data to databases or XML files

Session beans do not require database access, although they can obtain it indirectly as needed through entity beans. Session beans can also obtain direct access to databases (and other resources) through the use of resource references. Session beans can be either *stateful* or *stateless*.

Message-driven beans are new in version 2.0 of the Enterprise JavaBeans (EJB) specification. They enable asynchronous message servicing. The EJB container and a Java Message Service (JMS) provider work together to process messages. When a message arrives from another application component through

JMS, the EJB container forwards it through an `onMessage()` call to a message-driven bean instance, which then processes the message. In other respects, message-driven beans are similar to stateless session beans.

Beans that require data access use *data sources*, which are administrative resources that define pools of connections to persistent storage mechanisms.

For more information about enterprise beans, see "Resources for learning."

Developing enterprise beans

Design a J2EE application and the enterprise beans that it needs.

- For general design information, see "Resources for learning."
- Before developing entity beans with container-managed persistence (CMP), read "Concurrency control."

There are two basic approaches to selecting tools for developing enterprise beans:

- You can use one of the available integrated development environments (IDEs). IDE tools automatically generate significant parts of the enterprise bean code and contain integrated tools for packaging and testing enterprise beans. The IBM WebSphere Application Developer product is the recommended IDE. For more information, see the documentation for that product.
- If you have decided to develop enterprise beans without an IDE, you need at least an ASCII text editor. You can also use a Java development tool that does not support enterprise bean development. You can then use tools available in the Java Software Development Kit (SDK) and in this product to assemble, test, and deploy the beans.

The following steps primarily support the second approach, development without an IDE.

1. If necessary, migrate any pre-existing code to the required version of the Enterprise JavaBeans (EJB) specification.
2. Write and compile the components of the enterprise bean.
 - At a minimum, an EJB 1.1 session bean requires a bean class, a home interface, and a remote interface. An EJB 1.1 entity bean requires a bean class, a primary-key class, a home interface, and a remote interface.
 - At a minimum, an EJB 2.0 session bean requires a bean class, a home or local home interface, and a remote or local interface. An EJB 2.0 entity bean requires a bean class, a primary-key class, a remote home or local home interface, and a remote or local interface. The types of interfaces go together: If you implement a local interface, you must define a local home interface as well.

Note: Optionally, the primary-key class can be *unknown*. See unknown primary-key class for more information.

- Available only through EJB 2.0, a message-driven bean requires only a bean class.
3. For each entity bean, complete work to handle persistence operations.
 - Create a database schema for the entity bean's persistent data.
 - For entity beans with container-managed persistence (CMP), you must store the bean's persistent data in one of the supported databases. WebSphere Application Server application assembly tools automatically generate SQL code for creating database tables for CMP entity beans. If your CMP beans require complex database mappings, it is recommended that you use the IBM WebSphere Studio Application Developer product to generate code for the database tables.
 - For entity beans with bean-managed persistence (BMP), you can create the database and database table by using the database tools or use an existing database and database table.

For more information on creating databases and database tables, consult your database documentation.

- **(CMP entity beans for EJB 2.0 only)** Define finder queries with EJB Query Language (EJB QL).

With EJB QL, you define finders in terms of CMP fields and container-managed relationships, as follows:

- *Public* finders are visible in the bean's home interface. Implemented in the bean class, they return only remote interfaces and collection types.
- *Private* finders, expressed as SELECT statements, are used only within the bean class. They can return both local and remote interfaces, dependent values, other CMP field types, and collection types.
- **(CMP entity beans for EJB 1.1 only: an IBM extension)** Create a finder helper interface for each CMP entity bean that contains specialized finder methods (other than the `findByPrimaryKey` method).

The following logic is required for each finder method (other than the `findByPrimaryKey` method) contained in the home interface of an entity bean with CMP:

- The logic must be defined in a public interface named *NameBeanFinderHelper*, where *Name* is the name of the enterprise bean (for example, `AccountBeanFinderHelper`).
- The logic must be contained in a String constant named *findMethodName* WhereClause, where *findMethodName* is the name of the finder method. The String constant can contain zero or more question marks (?) that are replaced from left to right with the value of the finder method's arguments when that method is called.

5.1+ Assemble the beans in one or more EJB modules.

Migrating enterprise bean code to the supported specification

Support for Version 2.0 of the Enterprise JavaBeans (EJB) specification is new for Version 5 of this product. Migration of enterprise beans deployed in Version 4.0.x of this product is not generally necessary; Version 1.1 of the EJB specification is still supported. Follow these steps as appropriate for your application deployment.

1. Modify enterprise bean code for changes in the specification.
 - For Version 1.0 beans, migrate at least to Version 1.1.
 - As stated previously, migration from Version 1.1 to Version 2.0 of the EJB specification is not required for redeployment on this version of the product. However, if your application requires the capabilities of Version 2.0, migrate your Version 1.1-compliant code.

Note: The EJB Version 2.0 specification mandates that prior to the EJB container's executing a `findByMethod` query, the state of all enterprise beans enlisted in the current transaction be synchronized with the persistent store. (This is so the query is performed against current data.) If Version 1.1 beans are reassembled into an EJB 2.0-compliant module, the EJB container synchronizes the state of Version 1.1 beans as well as that of Version 2.0 beans. As a result, you might notice some change in application behavior even though the application code for the Version 1.1 beans has not been changed.

2. Modify enterprise bean code for changes in deployment requirements. If the enterprise beans were previously deployed in Version 3.0.x of this product, modify import statements to match standard package names. In Version 3.0.2.x, the following standard packages were present under nonstandard names:

```
javax.sql.*  
javax.transaction.*
```

Any code using WebSphere data sources, including BMP entity beans and session beans that access databases, must be modified.

3. You might have to modify code for some EJB 1.1-compliant modules that were not migrated to Version 2.0. Use the following information to help you decide.
 - Some stub classes for deployed enterprise beans have changed in the Java 2 Software Development Kit, Version 1.4.1.
 - The task of generating deployment code for enterprise beans changed significantly for EJB 1.1-compliant modules relative to EJB 1.0-compliant modules.

- If the CMP beans write to databases with mixed-case table or column names and you used IBM VisualAge for Java, Version 3.5.x, to generate the original deployment code, you cannot simply reassemble the beans in this product. You must export the original EJB project from the VisualAge for Java product as an EJB 1.1 JAR. This preserves the metadata needed to generate the correct deployment code for mixed-case database tables and columns. For more information, see the documentation for the Deployment Tool for Enterprise JavaBeans.

For detailed information about source and binary compatibility between deployed versions, see "Resources for learning."

4. Reassemble and redeploy all modules to incorporate migrated code.

Migrating enterprise bean code from Version 1.0 to Version 1.1

The following information generally applies to any enterprise bean that currently complies with Version 1.0 of the Enterprise JavaBeans (EJB) specification. For more information about migrating code for beans produced with the IBM WebSphere Studio Application Developer tool, see the documentation for that product. For more information about migrating code in general, see "Resources for learning."

1. In session beans, replace all uses of `javax.jts.UserTransaction` with `javax.transaction.UserTransaction`. Entity beans may no longer use the `UserTransaction` interface at all.
2. In finder methods for entity beans, include `FinderException` in the `throws` clause.
3. Remove `throws` of `java.rmi.RemoteException`; throw `javax.ejb.EJBException` instead. However, continue to include `RemoteException` in the `throws` clause of `home` and `remote` interfaces as required by the use of Remote Method Invocation (RMI).
4. Remove uses of the `finalize()` method.
5. Replace calls to `getCallerIdentity()` with calls to `getCallerPrincipal()`. The use of `getCallerIdentity()` is deprecated.
6. Replace calls to `isCallerInRole(Identity)` with calls to `isCallerInRole(String)`. The use of `isCallerInRole(Identity)` and `java.security.Identity` is deprecated.
7. Replace calls to `getEnvironment()` in favor of JNDI lookup. As an example, change the following code:

```
String homeName =
    aLink.getEntityContext().getEnvironment().getProperty("TARGET_HOME_NAME");
if (homeName == null) homeName = "TARGET_HOME_NAME";
```

The updated code would look something like the following:

```
Context env = (Context)(new InitialContext()).lookup("java:comp/env");
String homeName = (String)env.lookup("ejb10-properties/TARGET_HOME_NAME");
```

8. In `ejbCreate` methods for an entity bean with container-managed persistence (CMP), return the bean's primary key class instead of `void`.
9. Add the `getHomeHandle()` method to `home` interfaces.

```
public javax.ejb.HomeHandle getHomeHandle() {return null;}
```

Consider enhancements to match the following changes in the specification:

- Primary keys for entity beans can be of type `java.lang.String`.
- Finder methods for entity beans return `java.util.Collection`.
- Check the format of any JNDI names being used. Local name spaces are also supported.
- Security is defined by role, and isolation levels are defined at the method level rather than at the bean level.

Migrating enterprise bean code from Version 1.1 to Version 2.0

Enterprise JavaBeans (EJB) Version 2.0-compliant beans may be assembled only in an EJB 2.0-compliant module, although an EJB 2.0-compliant module can contain a mixture of Version 1.x and Version 2.0 beans.

The EJB Version 2.0 specification mandates that prior to the EJB container's executing a *findByMethod* query, the state of all enterprise beans enlisted in the current transaction be synchronized with the persistent store. (This is so the query is performed against current data.) If Version 1.1 beans are reassembled into an EJB 2.0-compliant module, the EJB container synchronizes the state of Version 1.1 beans as well as that of Version 2.0 beans. As a result, you might notice some change in application behavior even though the application code for the Version 1.1 beans has not been changed.

The following information generally applies to any enterprise bean that currently complies with Version 1.1 of the EJB specification. For more information about migrating code for beans produced with the IBM WebSphere Studio Application Developer tool, see the documentation for that product. For more information about migrating code in general, see "Resources for learning."

1. In beans with container-managed persistence (CMP) version 1.x, replace each CMP field with abstract get and set methods. In doing so, you must make each bean class abstract.
2. In beans with CMP version 1.x, change all occurrences of *this.field = value* to *setField(value)*.
3. In each CMP bean, create abstract get and set methods for the primary key.
4. In beans with CMP version 1.x, create an EJB Query Language statement for each finder method.
5. In finder methods for beans with CMP version 1.x, return *java.util.Collection* instead of *java.util.Enumeration*.
6. Update handling of non-application exceptions.
 - To report non-application exceptions, throw *javax.ejb.EJBException* instead of *java.rmi.RemoteException*.
 - Modify rollback behavior as needed: In EJB versions 1.1 and 2.0, all non-application exceptions thrown by the bean instance result in the rollback of the transaction in which the instance is running; the instance is discarded. In EJB 1.0, the container does not roll back the transaction or discard the instance if it throws *java.rmi.RemoteException*.
7. Update rollback behavior as the result of application exceptions.
 - In EJB versions 1.1 and 2.0, an application exception does not cause the EJB container to automatically roll back a transaction.
 - In EJB Version 1.1, the container performs the rollback only if the instance has called *setRollbackOnly()* on its *EJBContext* object.
 - In EJB Version 1.0, the container is required to roll back a transaction when an application exception is passed through a transaction boundary started by the container.

WebSphere extensions to the Enterprise JavaBeans specification

This article outlines extensions to the Enterprise JavaBeans (EJB) specification that IBM provides with this product:

Inheritance in enterprise beans

In the Java language, *inheritance* is the creation of a new class from an existing class or a new interface from an existing interface. This product supports two forms of inheritance: standard class inheritance and EJB inheritance.

In standard class inheritance, the home interface, remote interface, or enterprise bean class inherits properties and methods from base classes that are not themselves enterprise bean classes or interfaces.

By contrast in enterprise bean inheritance, an enterprise bean inherits properties (such as container-managed persistence (CMP) fields and container-managed relationship (CMR) fields), methods, and method-level control descriptor attributes from another enterprise bean.

For more information, see the documentation for the IBM WebSphere Studio Application Developer product.

Optimistic concurrency control for container-managed persistence

This product supports optimistic concurrency control of data access.

Access intents for EJB persistence

This product supports the application of named data-access policies.

Performance enhancements

Through the lifetime-in-cache settings, this product provides a way for you to improve performance for beans that are only occasionally updated. For more information, see "Entity bean assembly settings."

Some enterprise beans created with the IBM WebSphere Studio Application Developer product can utilize *read-ahead* for loading a bean and its related beans in a single database operation. An entire object graph or any part of the graph can be preloaded by configuring a finder method to use read-ahead.

Assembly and deployment extensions

5.1+ This product supports IBM extensions of assembly and deployment options.

Best practices for developing enterprise beans

Use the following guidelines when designing and developing enterprise beans:

- Use a stateless session bean to act as the entry point for business logic. For more information about using session facades, see "Resources for learning."
- Entity beans should use container-managed persistence.
- In an Enterprise JavaBeans (EJB) Version 2.0 environment, use local interfaces to improve communication between enterprise beans in the same Java virtual machine.
Local calls avoid the overhead of RMI/IIOP and use pass-by-reference semantics instead of pass-by-value. For each call, the caller and callee beans share the state of arguments. EJB 2.0 beans can have both a local and remote interface but more typically have one or the other.
- For communicating with remote clients, provide remote and remote home interfaces. For communicating with local clients like servlets, entity beans, and message-driven beans, provide local and local home interfaces.

Batch commands for container managed persistence

From JDBC 2.0 on, *PreparedStatement* objects can maintain a list of commands that can be submitted together as a batch. Instead of multiple database round trips, there can be only one database round trip for all the batched persistence requests.

The WebSphere Application Server version 5.0.2 enables you to take advantage of this. You can turn this option on from the EJB CMP side. When you choose this option, the run time defers *ejbStore/ejbCreate/ejbRemove* or the equivalent database persistence requests (insert/update/delete) until they are needed. This can be at the end of the transaction, or when a flush is needed for finders related to this EJB type. When the persistence operation finally happens, run time accumulates the database requests and uses JDBC *PreparedStatement* batch operation to make a single JDBC call for multiple rows of the same operation.

Setting the run time for batched commands:

1. Open the administrative console.
2. Select **Servers**.
3. Select **Application Servers**.
4. Select the server you want to configure.

5. In the Additional Properties area, select **Process Definition**.
6. In the Additional Properties area, select **Java Virtual Machine**.
7. Update the **Generic JVM arguments** with `Dcom.ibm.ws.pm.batch=true`.

Deferred Create for container managed persistence

The specification for Enterprise Java Beans (EJB) 2.x states that for Container Managed Persistence (CMP) during the `ejbCreate`, the container can create the representation of the entity in the database immediately, or defer it to a later time.

The WebSphere Application Server version 5.0.2 enables you to take advantage of this specification. You can turn this option on from the EJB CMP side. When you choose this option, the runtime defers `ejbCreate` (or the equivalent database persistence request) until it is needed. This can be at the end of the transaction, or when a flush is needed for finders related to this EJB type. By doing this you can reduce two round trips for the newly created entity (insert and update) to one (insert).

Setting the run time for deferred create:

1. Open the administrative console.
2. Select **Servers**.
3. Select **Application Servers**.
4. Select the server you want to configure.
5. In the Additional Properties area, select **Process Definition**.
6. In the Additional Properties area, select **Java Virtual Machine**.
7. Update the **Generic JVM arguments** with `Dcom.ibm.ws.pm.deferredcreate=true`.

Explicit invalidation in the Persistence Manager cache

Container managed persistence (CMP) entity beans can be configured as *long-lifetime* beans. A long-lifetime bean is one that is configured with *Lifetime In Cache Usage* equal to a value other than the default **OFF**. A value other than **OFF** means that data for this bean is cached beyond the end of the transaction in which the bean was obtained by a finder or other method. The *Lifetime In Cache Usage* and *Lifetime In Cache* values control the basic length of time the cached data remains valid. When the specified time runs out, the cached data is no longer valid. See the *LifetimeInCache* help sections of the Assembly Toolkit (ATK) for more details.

However, there is also an API that lets the client application code explicitly invalidate the cached data of a bean on demand, superceding the basic lifetime of the cache data as controlled by the *Lifetime In Cache Usage* and *Lifetime In Cache* settings. This is useful where an application that does not use CMP beans modifies the data that underlies a CMP bean (for example, it updates a database table to which a CMP bean is mapped). Such an application can inform WebSphere Application Server that any cached version of this bean data is **stale** and no longer matches what is in the database. The data should be invalidated (in essence, discarded). For CMP beans that cannot tolerate stale data, this is an important feature.

Because the PM Cache Invalidation mechanism does consume resources in exchange for its benefits, it is not enabled by default. To enable it refer to Setting Persistence Manager Cache Invalidation .

Example: Explicit Invalidation in the Persistence Manager Cache: Usage Scenario

The scenario of use for this feature begins with configuring one or more bean types to be long-lifetime beans (see Explicit Invalidation in the Persistence Manager Cache, and configuring the necessary Java Message Service (JMS) resources (described below). Once this is done, the server is started. The scenario continues as follows:

1. Assume that a CMP entity bean of type *Department* has been configured to be a long-lifetime bean.

2. Transaction 1 begins. Application code looks up *Department's* home and calls a finder method (such as *findByPrimaryKey("dept01")*). As this is the first finder to return *Department dept01*, a trip is made to the database to obtain the data. Transaction 1 ends.
3. Transaction 2 begins. Application code calls *findByPrimaryKey("dept01")* again. Because this is not the first finder to return *Department dept01*, we get a cache hit and no database trip is made. So far this is current WebSphere Application Server behavior for long-lifetime beans. Transaction 2 ends.
4. Another application, which does not use the *Department* CMP bean, is executed. This application might or might not be run on the WebSphere Application Server; it could be a legacy application. The application updates the database table that is mapped to the *Department* bean, altering the row for *dept01*. For example, the *budget* column in the table (mapped to a Java double CMP attribute in the *Department* bean) is changed from \$10,000.00 to \$50,000.00. This application was designed to cooperate with WebSphere Application Server. After performing the update, the application sends an invalidate request message to invalidate the *Department* bean *dept01*.
5. Transaction 3 begins. Application code looks up *Department's* home and calls a finder method (such as *findByPrimaryKey("dept01")*). Because this is the first finder after *Department dept01* is invalidated, a new database trip is made to obtain the data. Transaction 3 ends.

Persistence Manager cache invalidation API

The PM cache invalidation API is in the form of a JMS message that the client sends to a specially-named JMS topic using a connection from a specifically named JMS *TopicConnectionFactory*. The JMS message must be an *ObjectMessage* created by the client. The client code creates a *PMCacheInvalidationRequest* object that describes the bean data to invalidate. Client code places the *PMCacheInvalidationRequest* object in the *ObjectMessage* and publishes the *ObjectMessage* (for further details on the JMS objects and terms used here, please see the Java Message Service documentation).

The public class *PMCacheInvalidationRequest* is central to the API, so we include a portion of its code here for illustration purposes (if you see any differences between this illustration and the actual class, the class is to be considered correct):

```
package com.ibm.websphere.ejbpersistence;

/**
 *An instance of this class represents a request to invalidate one or more
 *CMP beans in the PMcache. When an invalidate occurs, cached data for this
 *bean is removed from the cache; the next time an application tries to find
 *this bean, a fresh copy of the bean data is obtained from the data store.
 *
 *The ability to invalidate a bean means that a CMP bean may be configured
 *as a long-lifetime bean and thus be cached across transactions for much
 *greater performance on future attempts to find this bean. Yet when some
 *outside mechanism updates the bean data, sending an invalidation request
 *will remove stale data from the PMcache so applications do not behave falsely
 *based on stale data.
 */
public class PMCacheInvalidationRequest implements Serializable {
    . . .

    /**
     * Constructor used to invalidate a single bean
     * @param beanHomeJNDIName the JNDI name of the bean home. This is the same value
     * used to look up the bean home prior to calling findByPrimaryKey, for example.
     * @param beanKey the primary key of the bean to be invalidated. The actual
     * object type must be the primary key type for this bean type.
     */
    public PMCacheInvalidationRequest(String beanHomeJNDIName, Object beanKey)
        throws IOException {
        . . .
    }

    /**
     * Constructor used to invalidate a Collection of beans
```



```

* @param beanHomeJNDIName java.lang.String the JNDI name of the bean home.
* This is the same value used to look up the bean home prior to calling
* findByPrimaryKey, for example.
* @param beanKeys a Collection of the primary keys of the beans to be
* invalidated. The actual type of each object in the Collection must be the
* primary key type for this bean type.
*/
public PMCacheInvalidationRequest(String beanHomeJNDIName, Collection beanKeys)
throws IOException {
    . . .
}
/**
* Constructor used to invalidate all beans of a given type
* @param beanHomeJNDIName java.lang.String the JNDI name of the bean home.
* This is the same value used to look up the bean home prior to calling
* findByPrimaryKey, for example.
*/
public PMCacheInvalidationRequest(String beanHomeJNDIName) {
    . . .
}
}

```

If the client wants to perform the invalidation in a synchronous way, it can opt to receive an acknowledgement JMS message when the invalidation is complete. To ask for an acknowledgement (ACK) message, the client sets a *Topic* of its own choosing in the *JMSReplyTo* field of the *ObjectMessage* for the invalidation request (see JMS documentation for further details). The client then waits (using the *receive()* method of JMS) on receipt of the acknowledgement message before continuing execution.

An ACK message enables the caller to insure there is not even a brief (seconds or less) window during which PM cache data is stale. The sending of an acknowledgement for each request does, of course, take a bit more time and so is recommended to be used only when needed.

The JMS resources used to make an invalidation request (*TopicConnectionFactory*, *TopicDestination*, and so forth) must be configured by the user (using the Administration console or other method) if they want to use PM Cache Invalidation. In this way the user can choose whichever JMS provider they prefer (as long as it supports pub-sub). The names that must be used for these resources are defined as part of the API, and use names unique to the WebSphere Application Server namespace to avoid name conflict with customer JMS resources.

The following are the names that must be used when the user configures the JMS resources (shown as Java constants for clarity):

```

// The JNDI name of the TopicConnectionFactory
private static final String topicConnectionFactoryJNDIName =
    "com.ibm.websphere.ejbpersistence.InvalidatetCF";
// The JNDI name of the TopicDestination
private static final String topicDestinationJNDIName = "com.ibm.websphere.ejbpersistence.invalidate";
// The Topic name (part of the TopicDestination)
private static final String topicString = "com.ibm.websphere.ejbpersistence.invalidate";

```

Here are examples of how these constants can be used in client code:

```

// Look up the TopicConnectionFactory...
InitialContext ic = new InitialContext();
TopicConnectionFactory topicConnectionFactory =
    (TopicConnectionFactory) ic.lookup(topicConnectionFactoryJNDIName);
...
// Look up the Topic
Topic topic = (Topic) ic.lookup(topicDestinationJNDIName);

```

Note that JMS messages can be sent not only from J2EE application code (for example, a *SessionBean* or *BMP* entity bean method) but also from non-J2EE applications if your chosen JMS provider allows for

this. For example, the IBM MQ provider, available in WebSphere Application Server as the **Embedded Messaging** feature (selectable during installation), supports the use of MQ classes (or structures in other languages) to create a topic connection and topic that are compatible with the *TopicConnectionFactory* and *TopicDestination* you configure using WebSphere Application Server Application Console.

Setting Persistence Manager Cache Invalidation:

1. Open the administrative console.
2. Select **Servers**.
3. Select **Application Servers**.
4. Select the server you want to configure.
5. In the Additional Properties area, select **Process Definition**.
6. In the Additional Properties area, select **Java Virtual Machine**.
7. Update the **Generic JVM arguments** with `-Dcom.ibm.ws.ejbpersistence.cacheinvalidation=true`.

Unknown primary-key class

When writing an entity bean for Enterprise Java Bean Version 2.0, the minimum requirements usually include a primary-key class. However, in some cases you might choose not to specify the primary-key class for an entity bean with container managed persistence (CMP). Perhaps there is no obvious primary key, or you want to allow the deployer to select the primary key fields at deployment time. The primary key type is usually derived from the type used by the database system that stores the entity objects, and you might not know what this key is.

So, the *unknown key type* is actually a type chosen at deployment time, making it changeable each time the bean is deployed. Your client code must deal with this key as type *Object*.

Currently, WebSphere Application Server supports top-down mapping and enables the deployer to choose *String* keys generated at the application server. For an example of how to use this function, see the Samples library.

Using access intent policies

You can use access intent policies to help the product run-time environment manage various aspects of Enterprise JavaBeans (EJB) persistence. You apply access intent policies to EJB Version 2.0 entity beans and their methods by using an application assembly tool. A set of default access intent policies comes with the Assembly Toolkit. You can also create your own custom policies.

1. Apply default access intent to CMP entity beans. For more information, see the online help available with the Assembly Toolkit.
2. Apply access intent policies to methods of CMP entity beans.
3. Create a custom access intent policy.
4. Apply access intent policies to BMP entity bean methods by using the *AccessIntent* API.
5. Apply multiple access intent policies to methods by using application profiling.

Access intent policies

An access intent policy is a named set of properties (access intents) that governs data access for Enterprise JavaBeans (EJB) persistence. You can assign policies to an entity bean and to individual methods on an entity bean's home, remote, or local interfaces during assembly. If you have the Integration Server product installed, you can assign these during development as well. You can set access intents only within EJB Version 2.x-compliant modules for entity beans with bean-managed persistence or with CMP Version 2.x.

This product supplies a number of access intent policies that specify permutations of read intent and concurrency control; the pessimistic/update policy can be qualified further. The selected policy determines the appropriate isolation level and locking strategy used by the run time environment.

Access intent policies are specifically designed to supplant the use of isolation level and access intent method-level modifiers found in the extended deployment descriptor for EJB version 1.1 enterprise beans. You cannot specify isolation level and read-only modifiers for EJB version 2.0 enterprise beans.

Access intent policies configured on an entity basis define the default access intent for that entity. The default access intent controls the entity unless you specify a different access intent policy based on either method-level configuration or application profiling

You can use application profiling or method level access intent policies to control access intent more precisely. Application profiling is only available in the Integration Server product. Method-level access intent policies are named and defined at the module level. A module can have one or many such policies. Policies are assigned, and apply, to individual methods of the declared interfaces of entity beans and their associated home interfaces. A policy is acted upon by either the combination of the EJB container and persistence manager (for entity beans with container-managed persistence) or directly by entity beans with bean-managed persistence.

For entity beans that are backed by tables with nullable columns, use an optimistic policy with caution. Nullable columns are automatically excluded from overqualified updates at deployment time; concurrent changes to a nullable field might result in lost updates. When used with the IBM WebSphere Studio Application Developer product, this product provides support for selecting a subset of the non-nullable columns that are to be reflected in the overqualified update statement that is generated in the deployment code to support optimistic policies.

An entity that is configured with a read-only policy that causes a bean to be activated can cause problems if updates are attempted within the same transaction. Those changes are not committed, and the process throws an exception because data integrity might be compromised.

Concurrency control

Concurrency control is the management of contention for data resources. A concurrency control scheme is considered *pessimistic* when it locks a given resource early in the data-access transaction and does not release it until the transaction is closed. A concurrency control scheme is considered *optimistic* when locks are acquired and released over a very short period of time at the end of a transaction.

The objective of optimistic concurrency is to minimize the time over which a given resource would be unavailable for use by other transactions. This is especially important with long-running transactions, which under a pessimistic scheme would lock up a resource for unacceptably long periods of time.

Under an optimistic scheme, locks are obtained immediately before a read operation and released immediately afterwards. Update locks are obtained immediately before an update operation and held until the end of the transaction.

To enable optimistic concurrency, this product uses an *overqualified update scheme* to test whether the underlying data source has been updated by another transaction since the beginning of the current transaction. With this scheme, the columns marked for update and their original values are added explicitly through a WHERE clause in the UPDATE statement so that the statement fails if the underlying column values have been changed. As a result, this scheme can provide column-level concurrency control; pessimistic schemes can control concurrency at the row level only.

Optimistic schemes typically perform this type of test only at the end of a transaction. If the underlying columns have not been updated since the beginning of the transaction, pending updates to container-managed persistence fields are committed and the locks are released. If locks cannot be

acquired or if some other transaction has updated the columns since the beginning of the current transaction, the transaction is rolled back: All work performed within the transaction is lost.

Pessimistic and optimistic concurrency schemes require different transaction isolation levels. Enterprise beans that participate in the same transaction and require different concurrency control schemes cannot operate on the same underlying data connection.

Whether or not to use optimistic concurrency depends on the type of transaction. Transactions with a high penalty for failure might be better managed with a pessimistic scheme. (A high-penalty transaction is one for which recovery would be risky or resource-intensive.) For low-penalty transactions, it is often worth the risk of failure to gain efficiency through the use of an optimistic scheme. In general, optimistic concurrency is more efficient when update collisions are expected to be infrequent; pessimistic concurrency is more efficient when update collisions are expected to occur often.

Read-ahead hints

Read-ahead schemes enable applications to minimize the number of database roundtrips by retrieving a working set of container-managed persistence (CMP) beans for the transaction within one query.

Read-ahead involves activating the requested CMP beans and caching the data for their related beans, which ensures that data is present for the beans that are most likely to be needed next by an application.

A *read-ahead hint* is a canonical representation of the related beans that are to be read. It is associated with the *findByPrimaryKey* method for the requested bean type, which must be an EJB 2.x-compliant CMP entity bean.

Read-ahead hints can be set only using the WebSphere Business Integration Server Foundation assembly tool or through the Add Access Intent wizard of the IBM WebSphere Studio Application Developer product.

Read-ahead is only supported for access intent policies that can be applied by the backend against which the application is deployed. Otherwise, the read-ahead hint is disregarded.

Currently, only *findByPrimaryKey* methods can have read-ahead hints. Only beans related to the requested beans by a container-managed relationship (CMR), either directly or indirectly through other beans, can be read ahead. Beans that use EJB inheritance should not be used in a read-ahead hint.

A read-ahead hint takes the form of a character string. You do not have to provide the string; the wizard generates it for you based on CMRs defined for the bean. The following example is provided as supplemental information only.

Suppose a CMP bean type A has a finder method that returns instances of bean A. A read-ahead hint for this method is specified using the following notation: *RelB.RelC; RelD*

Interpret the preceding notation as follows:

- Bean type A has a CMR with bean types B and D.
- Bean type B has a CMR with bean type C.

For each bean of type A that is retrieved from the database, its directly-related B and D beans and its indirectly-related C beans are also retrieved. The order of the retrieved bean data columns in each row of the result set is the same as their order in the read-ahead hint: an A bean, a B bean (or null), a C bean (or null), a D bean (or null). For hints in which the same relationship is mentioned more than once (for example, *RelB.RelC;RelB.RelE*), a bean's data columns appear only once, at the position it first appears in the hint.

The tokens shown in the notation (*RelB* and so on) must be CMR field names for the relationships as defined in the deployment descriptor for the bean. In indirect relationships such as *RelB.RelC*, *RelC* is a CMR field name defined in the deployment descriptor for bean type B.

A single read-ahead hint cannot refer to the same bean type in more than one relationship. For example, if a Department bean has a relationship *employees* with the Employee bean and also has a relationship *manager* with the Employee bean, the read-ahead hint cannot specify both *employees* and *manager*.

For more information about how to set read-ahead hints, see the documentation for the Websphere Studio Application Developer product.

Access intent service

Access intent is a WebSphere Application Server run-time service that enables you to more precisely manage an application's persistence. The access intent service defines a set of declarative annotations used by the Enterprise JavaBeans (EJB) container and its agents to make performance optimizations for entity bean access. These annotations are organized into sets called *access intent policies*.

Access intent policies contain a set of annotations considered as hints by the EJB container and its agents. Most access intent policies are hints representing high-level abstractions that can be mapped to a specific backend resource manager. It is the responsibility of the EJB persistence machinery to ensure the necessary concurrency control, connection, and cache management when carrying out the persistence details. The EJB persistence manager can use access intent hints to make better performance decisions when carrying out its assigned task. A smaller number of access intents are hints to the EJB container, influencing the management of EJB collections.

Although it is recommended that you always configure bean level access intent for your applications, if you find it necessary you can apply access intent policies to methods within the scope of an EJB module. In such cases the policy becomes the default access intent for all requests upon the configured methods.

You can also apply access intent policies to beans within the scope of application profiles. Consequently, you can configure beans with multiple and opposing access intent policies. The application profiling documentation explains in more detail how to configure an application to apply a particular access intent policy to a bean for one request, then apply another access intent policy to the same bean for a different request.

Access intent with BMP entity beans

Access intent's declarative functionality provides great power to you as a CMP entity bean developer. You can provide hints on how WebSphere Application Server is to manage the details of persistence without having to explicitly manage any of the persistence logic from within the application.

There are situations, however, in which you might need to develop BMP entity beans. Because the only meaningful difference between BMP and CMP components is who provides the persistence logic, BMP entity beans should be able to leverage access intent hints just as WebSphere Application Server does on behalf of CMP entity beans. BMP entity beans that use the access intent service participate in application profiling; that is, the value of the access intent attributes can differ from request to request, allowing the BMP entity bean to seamlessly modify its persistence strategy.

You can apply access intent policies to BMP entity bean methods as well as CMP entity bean methods. Because access intent hints are not contractual in nature, there is no obligation for a BMP entity bean to exploit them. BMP entity beans are expected to use only those access intent attributes that are important to that particular bean.

The current access intent policy is bound into the `java:comp` namespace for a particular BMP entity bean. That policy is current only for the duration of the method call during which the access intent policy was retrieved. In a typical scenario, you would cache the access type during invocation of the `ejbLoad()` method so that appropriate actions can be taken during invocation of the `ejbStore()` method.

Access intent design considerations

Use the access intent service to solve clear performance problems. Identify usage patterns that lead to poor application performance and apply appropriate access intent policies.

Refrain from over-tuning an application. You can introduce errors by incorrectly using the access intent service. For example, misuse of the `wsPessimisticUpdate-NoCollision` policy can result in lost updates; inappropriately setting the collection increment value can introduce performance issues; and problem determination is more difficult when an application is confusingly configured with multiple access intent policies. Clarity and simplicity should be your guiding principles when using the access intent service. This is even more important when applying access intent policies within the scope of application profiles (a feature of WebSphere Business Integration Server Foundation).

Even though access intent policies can be configured on any method of an entity bean, some attributes of a policy can only be leveraged by the run-time environment under certain conditions. For example, concurrency and access intent are only used for CMP entity beans when the `ejbLoad()` method is driven to open a connection to, and read data from, a given resource; that data is cached and used to drive the proper queries during invocation of the `ejbStore()` method. Read-ahead hints are only used during the execution of a finder for a bean. Finally, the collection increment and resource manager prefetch increment are only used on multi-object finders. Configuring policies on methods that will not use the policy is not an error (only certain attributes of any policy are used, even when the policy is appropriately applied to a method). However, configuring policies unnecessarily throughout an application obscures the design of the application and complicates the maintenance of the application.

Applying access intent policies to methods

You apply an access intent policy to a method, or set of methods, in an application's entity beans through the Assembly Toolkit.

1. Start the Assembly Toolkit.
2. Optional: Open the J2EE perspective to work with J2EE projects. Click **Window > Open Perspective > Other > J2EE**.
3. Optional: Open the J2EE Hierarchy view. Click **Window > Show View > J2EE Hierarchy**. Other helpful views include the Project Navigator view (**Window > Show View > Other > J2EE > Project Navigator**) and the Navigator view (**Window > Show View > Navigator**).
4. Select the **Access** tab.
5. On the right side of the **Access Intent for Entities 2.x (Method Level)** panel, select **Add**. The **Add Access Intent** panel displays.
6. Specify the **Name** for your new intent policy.
7. Select the **Access intent name** from the drop-down list.
8. Enter a **Description** to help you remember what this policy does.
9. Optional: Select **Read Ahead Hint**.
10. Click **Next**. The next **Add Access Intent** panel displays, with optional attributes.
11. Optional: Decide whether or not to overwrite these optional access intent attributes. Click on those you want to change.
12. Click **Next**. The next **Add Access Intent** panel, with a list of Enterprise Beans, displays.
13. Select one or more Enterprise Beans from the list.

Note: If you selected **Read Ahead Hint** in an earlier step, you can only select **ONE** bean at this step.

14. Click **Next**. The next **Add Access Intent** panel, with a list of methods, displays.
15. Select the methods you want to use.

16. If you *DID NOT* select **Read Ahead Hint** in an earlier step, click **Finish**. If you *DID* select the Read Ahead Hint option, you can click **Next** to specify your Read Ahead Hint for the specified bean. The next **Add Access Intent** panel, with a list of EJB preload paths, displays.
17. Edit the EJB preload path by selecting relationship roles from the **Relationship roles:** window.
18. Click **Finish**. A new entry is created in the **Access Intent for Entities 2.x (Method Level)** panel

Using the AccessIntent API

This task describes how to programmatically retrieve and call the AccessIntent API during the execution of BMP entity bean methods.

1. Look up the current access intent in the namespace. For example:

```
InitialContext ic = new InitialContext();
AccessIntent ai = ic.lookup("java:comp/websphere/AppProfile/AccessIntent");
```

2. Call the necessary get() methods. For example:

```
int concurrency = ai.getConcurrencyControl();
int accessType = ai.getAccessType();
if ( (concurrency == AccessIntent.CONCURRENCY_CONTROL_PESSIMISTIC)
    && (accessType == AccessIntent.ACCESS_TYPE_UPDATE) ) {
    boolean exclusive = ai.getPessimisticUpdateHintExclusive();
    // . . .
}
// . . .
```

Note: The access intent object reference retrieved from the java:comp lookup is current for the duration of the method in which the reference was looked up. Depending on how you configured the application profile, subsequent calls of the same method might not retrieve the same access intent reference. You can only look up the object reference during the call of a BMP entity bean's method; the reference does not exist during a request on a CMP entity bean. Therefore, access intent object references should not be cached beyond, or used outside of, the scope of the execution of any given BMP method.

AccessIntent interface

The AccessIntent interface is available to BMP entity beans by the following JNDI lookup:

```
java:comp/websphere/AppProfile/AccessIntent
```

AccessIntent interface

```
package com.ibm.websphere.appprofile.accessintent;

/**
 * This interface defines the essential access intents
 * available at runtime.
 */
public interface AccessIntent {

    /**
     * Returns the concurrency control intent, which indicates
     * the application prefers either pessimistic or optimistic
     * concurrency control when accessing the current component
     * in the context of the current transaction.
     */
    public int getConcurrencyControl();
    public final int CONCURRENCY_CONTROL_PESSIMISTIC = 1;
    public final int CONCURRENCY_CONTROL_OPTIMISTIC = 2;

    /**
     * Returns access type intent, which indicates the application
     * intends either update or read access of the current component
     * in the context of the current transaction.
     */
}
```



```

public int getAccessType();
public final int ACCESS_TYPE_UPDATE= 1;
public final int ACCESS_TYPE_READ = 2;

/**
 * Returns a boolean where true indicates that the runtime should
 * assume that there will be no collision on retrieved rows.
 */
public boolean getPessimisticUpdateHintNoCollision();

/*
 * Returns a boolean where true indicates that the runtime should
 * assume that there will be collisions on retrieved rows.
 */
public boolean getPessimisticUpdateHintExclusive();

/**
 * Returns the collection access intent, which indicates the
 * application intends to access the objects returned by the
 * currently executing finder in either serial or random fashion.
 */
public int getCollectionAccess();
public final int COLLECTION_ACCESS_RANDOM = 1;
public final int COLLECTION_ACCESS_SERIAL = 2;

/**
 * Returns the collection scope, which indicates the maximum
 * lifespan of a lazy collection.
 */
public int getCollectionScope();
public final int COLLECTION_SCOPE_TRANSACTION = 1;
public final int COLLECTION_SCOPE_ACTIVITYSESSION = 2;
public final int COLLECTION_SCOPE_TIMEOUT = 3;

/**
 * Returns the timeout value in seconds when collectionScope is Timeout.
 */
public int getCollectionTimeout();

/**
 * Returns the number of elements the application requests be contained
 * in each segment of the element collection returned by the currently
 * executing finder.
 */
public int getCollectionIncrement();

/**
 * Returns the ReadAheadHint requested by the application for the currently
 * executing finder.
 */
public ReadAheadHint getReadAheadHint();

/**
 * Returns the number of elements the application requests be contained in
 * each segment of a a query made on a database.
 */
public int getResourceManagerPreFetchIncrement();
}

```

Access intent exceptions

The following exceptions are thrown in response to the application of access intent policies:

com.ibm.ws.ejbpersistence.utilpm.PersistenceManagerException

If the method that drives the `ejbLoad()` method is configured to be read-only but updates are then made within the transaction that loaded the bean's state, an exception is thrown during invocation

of the `ejbStore()` method, and the transaction is rolled back. Likewise, the `ejbRemove()` method cannot succeed in a transaction that is set as read-only. If an update hint is applied to methods of entity beans with bean-managed persistence, the same behavior and exception results. The forwarded exception object contains the message string `PMGR1103E: update instance level read only bean beanName`

This exception is also thrown if the applied access intent policy cannot be honored because a finder, `ejbSelect`, or container-managed relationship (CMR) accessor method returns an inherently read-only result. The forwarded exception object contains the message string `PMGR1001: No such DataAccessSpec - methodName`

The most common occurrence of this error is when a custom finder that contains a read-only EJB Query Language (EJB QL) statement is called with an applied access intent of `wsPessimisticUpdate` or `wsPessimisticUpdate-Exclusive`. These policies require the use of a FOR UPDATE clause on the SQL SELECT statement to be executed, but a read-only query cannot support FOR UPDATE. Other examples of read-only queries include joins; the use of ORDER BY, GROUP BY, and DISTINCT keywords.

To eliminate the exception, edit the EJB query so that it does not return an inherently read-only result or change the access intent policy being applied.

- If an update access is required, change the applied access intent setting to `wsPessimisticUpdate-WeakestLockAtLoad` or `wsOptimisticUpdate`.
- If update access is not truly required, use `wsPessimisticRead` or `wsOptimisticRead`.
- If connection sharing between entity beans is required, use `wsPessimisticUpdate-WeakestLockAtLoad` or `wsPessimisticRead`.

com.ibm.websphere.ejb.container.CollectionCannotBeFurtherAccessed

If a lazy collection is driven after it is no longer in scope, and beyond what has already been locally buffered, a `CollectionCannotBeFurtherAccessed` exception is thrown.

com.ibm.ws.exception.RuntimeWarning

If an application is configured incorrectly, a run-time warning exception is thrown as the application starts; startup is ended. You can validate an application's configuration by choosing the `verify` function. Some examples of misconfiguration include:

- A method configured with two different access intent policies
- A method configured with an undefined access intent policy

javax.ejb.NoSuchEntityException

If an update fails under optimistic concurrency because fields changed within another transaction between load and store requests, a `NoSuchEntityException` is raised and the commit fails.

Access intent assembly settings

Access intent policies contain data-access settings for use by the persistence manager. Default access intent policies are configured on the entity bean. Optionally, you can associate access intent policies with one or more methods.

These settings are applicable only for EJB 2.x-compliant entity beans that are packaged in EJB 2.x-compliant modules. Connection sharing between beans with bean-managed persistence and those with container-managed persistence is possible if they all use the same access intent policy.

Name

Specifies a name for a mapping between an access intent policy and one or more methods.

Description

Contains text that describes the mapping.

Methods - Name

Specifies the name of an enterprise bean method, or the asterisk character (*). The asterisk is used to denote all of the methods of an enterprise bean's remote and home interfaces.

Methods - Enterprise bean

Specifies which enterprise bean contains the methods indicated in the Name setting.

Methods - Type

Used to distinguish between a method with the same signature that is defined in both the home and remote interface. Use `Unspecified` if an access intent policy applies to all methods of the bean.

Data type	String
Range	Valid values are Home, Remote, Local, LocalHome or Unspecified

Methods - Parameters

Contains a list of fully qualified Java type names of the method parameters. This setting is used to identify a single method among multiple methods with an overloaded method name.

Applied access intent

Specifies how the container must manage data access for persistence. Configurable both as a default access intent for an entity and as part of a method-level access intent policy.

Data type	String
Default	<code>wsPessimisticUpdate-WeakestLockAtLoad</code> . With Oracle, this is the same as <code>wsPessimisticUpdate</code> .
Range	Valid settings are <code>wsPessimisticUpdate</code> , <code>wsPessimisticUpdate-NoCollision</code> , <code>wsPessimisticUpdate-Exclusive</code> , <code>wsPessimisticUpdate-WeakestLockAtLoad</code> , <code>wsPessimisticRead</code> , <code>wsOptimisticUpdate</code> , or <code>wsOptimisticRead</code> . Only <code>wsPessimisticRead</code> and <code>wsOptimisticRead</code> are valid when class-level caching is enabled in the EJB container.

This product supports lazy collections. For each segment of a collection, iterating through the collection (`next()`) does not trigger a remote method call to retrieve the next remote reference. Two policies (`wsPessimisticUpdate` and `wsPessimisticUpdate-Exclusive`) are extremely lazy; the collection increment size is set to 1 to avoid overlocking the application. The other policies have a collection increment size of 25.

If an entity is not configured with an access intent policy, the run-time environment typically uses `wsPessimisticUpdate-WeakestLockAtLoad` by default. If, however, the **Lifetime in cache** property is set on the bean, the default value of **Applied access intent** is `wsOptimisticRead`; updates are not permitted.

Additional information about valid settings follows:

Profile name	Concurrency control	Access type	Transaction isolation
<code>wsPessimisticRead</code> (Note 1)	pessimistic	read	For Oracle, read committed. Otherwise, repeatable read
<code>wsPessimisticUpdate</code> (Note 2)	pessimistic	update	For Oracle, read committed. Otherwise, repeatable read
<code>wsPessimisticUpdate-Exclusive</code> (Note 3)	pessimistic	update	serializable
<code>wsPessimisticUpdate-NoCollision</code> (Note 4)	pessimistic	update	read committed
<code>wsPessimisticUpdate-WeakestLockAtLoad</code> (Note 5)	pessimistic	update	Repeatable read

Profile name	Concurrency control	Access type	Transaction isolation
wsOptimisticRead	optimistic	read	read committed
wsOptimisticUpdate (Note 6)	optimistic	update	read committed

Notes:

1. Read locks are held for the duration of the transaction.
2. The generated SELECT FOR UPDATE query grabs locks at the beginning of the transaction.
3. SELECT FOR UPDATE is generated; locks are held for the duration of the transaction.
4. A plain SELECT query is generated. No locks are held, but updates are permitted. Use cautiously. This intent enables execution without concurrency control.
5. Where supported by the backend, the generated SELECT query does not include FOR UPDATE; locks are escalated by the persistent store at storage time if updates were made. Otherwise, the same as wsPessimisticUpdate.
6. Generated overqualified-update query forces failure if CMP column values have changed since the beginning of the transaction.

Be sure to review the rules for forming overqualified-update query predicates. Certain column types (for example, BLOB) are ineligible for inclusion in the overqualified-update query predicate and might affect your design.

Access intent best practices

This topic outlines issues to consider when applying access intent policies to Enterprise JavaBeans (EJB) methods.

- **Start by configuring the default access intent policy for an entity.** After your application is built and running, you can more finely tune certain access paths in your application using application profiling or method-level access intent.
- **Don't mix access types.** Avoid using both pessimistic and optimistic policies in the same transaction. For most databases, pessimistic and optimistic policies use different isolation levels. This can result in multiple database connections, which prevents you from taking advantage of the performance benefits possible through connection sharing.
- **Take care when applying wsPessimisticUpdate-NoCollision.** This policy does not ensure data integrity. No database locks are held, so concurrent transactions can overwrite each other's updates. Use this policy only if you can be sure that only one transaction will attempt to update persistent store at any given time.

Frequently asked questions: Access intent

I have not applied any access intent policies at all. My application runs just fine with a DB2 database, but it fails with an Oracle database with the following message:

```
com.ibm.ws.ejbpersistence.utilpm.PersistenceManagerException: PMGR1001E: No such DataAccessSpec:FindAllCustomers. The backend datastore does not support the SQLStatement needed by this AccessIntent: (pessimistic update-weakestLockAtLoad)(collections: transaction/25) (resource manager prefetch: 0) (AccessIntentImpl@d23690a). Why?
```

If you have not configured access intent, all of your data is accessed under the default access intent policy (wsPessimisticUpdate-WeakestLockAtLoad). On DB2 databases, the weakest lock is a shared one, and the query runs without a FOR UPDATE clause. On Oracle databases, however, the weakest lock is an update lock; this means that the SQL query must contain a FOR UPDATE clause. However, not every SQL statement necessarily supports FOR UPDATE; for example, if the query is being run against multiple tables in a join, FOR UPDATE is not supported.

To avoid this problem, try either of the following:

- Modify your SQL query or reconfigure your application so that an update lock is supported
- Apply an access intent policy that supports optimistic concurrency

I am calling a finder method and I get an InconsistentAccessIntentException at run time. Why?

This can occur when you use method-level access intent policies to apply more control over how a bean instance is loaded. This exception indicates that the entity bean was previously loaded in the same transaction. This could happen if you called a multifinder method that returned the bean

instance with access intent policy X applied; you are now trying to load the second bean again by calling its `findByPrimaryKey` method with access intent Y applied. Both methods must have the same access intent policy applied.

Likewise, if the entity was loaded once in the transaction using an access intent policy configured on a finder, you might have called a container-managed relationship (CMR) accessor method that returned the entity bean configured to load using that entity's default access intent.

To avoid this problem, ensure that your code does not load the same bean instance twice within the same transaction with different access intent policies applied. Avoid the use of method-level access intent unless absolutely necessary.

I have two beans in a container-managed relationship. I call `findByPrimaryKey()` on the first bean and then call `getBean2()`, a CMR accessor method, on the returned instance. At that point, I get an `InconsistentAccessIntentException`. Why?

You are probably using read-ahead. When you loaded the first bean, you caused the second bean to be loaded under the access intent policy applied to the finder method for the first bean.

However, you have configured your CMR accessor method from the first bean to the second with a different access intent policy. CMR accessor methods are really finder methods in disguise; the run-time environment behaves as if you were trying to change the access intent for an instance you have already read from persistent store.

To avoid this problem, beans configured in a read-ahead hint are all driven to load with the same access intent policy as the bean to which the read-ahead hint is applied.

I have a bean with a one-to-many relationship to a second bean. The first bean has a pessimistic-update intent policy applied. When I try to add an instance of the second bean to the first bean's collection, I get an `UpdateCannotProceedWithIntegrityException`. Why?

The second bean probably has a read intent policy applied. When you add the second bean to the first bean's collection, you are not updating the first bean's state, you are implicitly modifying the second bean's state. (The second bean contains a foreign key to the first bean, which is modified.)

To avoid this problem, ensure that both ends of the relationship have an update intent policy applied if you expect to change the relationship at run time.

EJB modules

An EJB module is used to assemble one or more enterprise beans into a single deployable unit. An EJB module is stored in a standard Java archive (JAR) file.

An EJB module contains the following:

- One or more deployable enterprise beans.
- A deployment descriptor, stored in an Extensible Markup Language (XML) file. This file declares the contents of the module, defines the structure and external dependencies of the beans in the module, and describes how the beans are to be used at run time.

You can deploy an EJB module as a stand alone application, or combine it with other EJB modules or with Web modules to create a J2EE application. An EJB module is installed and run in an enterprise bean container.

For more information about EJB modules, see "Resources for learning."

Assembling EJB modules

Assemble an Enterprise JavaBeans (EJB) module to contain enterprise beans and related code artifacts. Group Web components, client code, and resource adapter code in separate modules. After assembling an EJB module, you can install it as a stand-alone application or combine it with other modules into an enterprise application.

To increase performance, break container-managed persistence (CMP) enterprise beans into several enterprise bean modules during assembly. The load time for hundreds of beans is improved by distributing the beans across several JAR files and packaging them to an EAR file. Load time is faster when the administrative server attempts to start the beans, for example, 8-10 minutes versus more than one hour when one JAR file is used.

Use the Assembly Toolkit to assemble an EJB module in any of the following ways:

- Import an existing EJB module (EJB JAR file).
 - Create a new EJB module.
 - Copy code artifacts (such as entity beans) from one EJB module into a new EJB module.
1. Start the Assembly Toolkit.
 2. Optional: Open the J2EE perspective to work with J2EE projects. Click **Window > Open Perspective > Other > J2EE**.
 3. Optional: Open the J2EE Hierarchy view. Click **Window > Show View > J2EE Hierarchy**. Other helpful views include the Project Navigator view (**Window > Show View > Other > J2EE > Project Navigator**) and the Navigator view (**Window > Show View > Navigator**).
 4. **5.1+** Optional: To use an application assembly service of WebSphere Business Integration (WBI) Server Foundation, select the **Integration Server** target server. Available WBI assembly services include ActivitySession, Application profiling, Extended access, Extended messaging, Internationalization, Last participant support, and Pushdown.
 - a. Click **Project > Properties > J2EE**.
 - b. For **Target server**, select **Integration Server**. For this release, **Integration Server v5.1**.
 - c. Click **OK**.
 5. Migrate enterprise bean (JAR) files created with the Application Assembly Tool (AAT) or a different tool to the Assembly Toolkit. To migrate files, import your enterprise bean files to the Assembly Toolkit.
 6. Create a new EJB module.
 7. Copy code artifacts (such as entity beans) from one EJB module into a new EJB module.
 8. Verify the contents of the new EJB module in either of the following ways:
 - In the J2EE Hierarchy view, expand **EJB Modules** and view the new module.
 - Click **Window > Show View > Navigator** to see the associated files for the EJB module in a Navigator view.

Container transactions

Container transaction properties specify how an EJB container is to manage transaction scopes for the enterprise bean's method invocations. A transaction attribute is mapped to one or more methods.

Method extensions

Method extensions are IBM extensions to the standard deployment descriptors for enterprise beans.

Method extension properties are used to define transaction isolation levels for methods, to control the delegation of a principal's credentials, and to define custom finder methods.

Method permissions

A method permission is a mapping between one or more security roles and one or more methods that a member of the role can call.

References

References are logical names used to locate external resources for enterprise applications. References are defined in the application's deployment descriptor file. At deployment, the references are bound to the physical location (global JNDI name) of the resource in the target operational environment.

This product supports the following types of references:

- An EJB reference is a logical name used to locate the home interface of an enterprise bean.
- A resource reference is a logical name used to locate a connection factory object.

These objects define connections to external resources such as databases and messaging systems. The container makes references available in a JNDI naming subcontext. By convention, references are organized as follows:

- EJB references are made available in the `java:comp/env/ejb` subcontext.
- Resource references are made available as follows:
 - JDBC DataSource references are declared in the `java:comp/env/jdbc` subcontext.
 - JMS connection factories are declared in the `java:comp/env/jms` subcontext.
 - JavaMail connection factories are declared in the `java:comp/env/mail` subcontext.
 - URL connection factories are declared in the `java:comp/env/url` subcontext.

EJB containers

An Enterprise JavaBeans (EJB) container provides a run-time environment for enterprise beans within the application server. The container handles all aspects of an enterprise bean's operation within the application server and acts as an intermediary between the user-written business logic within the bean and the rest of the application server environment.

One or more EJB modules, each containing one or more enterprise beans, can be installed in a single container.

The EJB container provides many services to the enterprise bean, including the following:

- Beginning, committing, and rolling back transactions as necessary.
- Maintaining pools of enterprise bean instances ready for incoming requests and moving these instances between the inactive pools and an active state, ensuring that threading conditions within the bean are satisfied.
- Most importantly, automatically synchronizing data in an entity bean's instance variables with corresponding data items stored in persistent storage.

By dynamically maintaining a set of active bean instances and synchronizing bean state with persistent storage when beans are moved into and out of active state, the container makes it possible for an application to manage many more bean instances than could otherwise simultaneously be held in the application server's memory. In this respect, an EJB container provides services similar to virtual memory within an operating system.

Between transactions, the state of an entity bean can be cached. The EJB container supports option A, B, and C caching.

For more information about EJB containers, see "Resources for learning."

Container Managed Persistence over anything

Before the release of WebSphere Business Integration Server Foundation Version 5.1, Container Managed Persistence (CMP) supported a number of *relational* backend systems. Users wanting to use WebSphere Application Server Enterprise Edition to host stateful and live business models needed a way to represent all the entities that exist in the business model, including *non-relational* backend systems. Version 5.1 addresses this need.

The CMP infrastructure is designed to work with the J2EE Connector Architecture (JCA), through an interface called *EJBToRAAdapter*. In addition to the current implementation of this interface (which connects to *relational* backend systems), Version 5.1 provides a new implementation that connects to *procedural* backend systems. The EJB Deployment tool can be used to return either of these

implementations. The new procedural implementation contains generic logic for adapting between JCA and the WebSphere Application Server persistence manager. Note that all of the actual connection logic is contained down in the *FunctionSet* of the bean.

However, the CMP over anything (CMP/A) function is not limited to JCA connectors. There are other approaches to backend data store connectivity that you can use. CMP over the web services, which enable communication with backend systems indirectly through a web service using JAX-RPC, or WSIF are available. You can also use a composed EJB model, where a CMP bean delegates its persistence logic to another EJB.

So now you can develop a simple CMP bean that works with whatever backend system you need, rather than having to go the cumbersome Bean Managed Persistence (BMP) bean development route, where all of the connection logic must go directly in the bean implementation class.

Version 5.1 introduces support for *pushdown* methods, which can be thought of as *container managed business methods*. In the CMP framework, you can use the class called *UserDefinedPushDownMethodImpl*, to place connectivity logic with a non-relational backend data store. The Java source code for this class is generated by the WebSphere Application Server deployment tools and contains methods corresponding to each CRUD method and each pushdown data-logic method defined in the PME EJB deployment descriptor.

This means that methods can be defined on a CMP bean that, rather than being implemented by the developer (in a BMP style of putting connection logic directly in the EJB's method implementation), are instead handled by the container and deployment tools. This is accomplished by defining the method as **abstract**, and then marking new PME extended deployment descriptor settings for these *PushDownMethodElements*. The new command-line code generation tool reads these extensions and generates appropriate entries.

WebSphere Business Integration Server Foundation Version 5.1 also includes a generic *UserDefinedPushDownMethods* interface (unrelated to the backend) and a *UserDefinedPushDownMethodImpl* class that implements the interface and that is specific to a given backend. When a method has been marked **userDefined**, the *FunctionSet* is delegated to the corresponding method of the *UserDefinedPushDownMethodImpl*. The implementer of the *UserDefinedPushDownMethodImpl* does not need to worry about processing CCI Records - it can simply do its work and let the *FunctionSet* code take care of that.

See Example: Container Managed Persistence over anything -- CCI and Example: Container Managed Persistence over anything -- EJB for examples.

Version 5.1 provides tooling that you can invoke from the command line or from the Assembly Toolkit. The tool generates the bean's *UserDefinedPushDownMethods* interface, and generates skeleton Java source code for the bean's *UserDefinedPushDownMethodImpl* class. You must fill this class in with the appropriate logic, compile it, and package it into the appropriate EJB JAR file. See Applying Container Managed Persistence services over anything for complete information.

The following links provide additional information concerning the procedural functions that use a CMP EJB. For relational stored procedure information, refer to "Stored procedure support for CMP EJB" in the Information Center.

Applying Container Managed Persistence services over anything

1. Develop your Container Managed Persistence (CMP) bean as usual. See Developing enterprise beans for more information. Optionally, you can also define abstract methods that are treated as pushdown methods.
2. Assemble the Enterprise JavaBean (EJB) Java Archive (JAR) file. In addition to standard packaging steps, make sure to set the proper deployment descriptor settings.

- a. Use the Assembly Toolkit to make your settings.
- b. There is a *PushDownContainerManagedEntityExtension* method associated with your CMP bean, and its *procedural* attribute is set to **True**. Set its *backEndType* attribute to one of the following:

CCI Creates skeleton for the *UserDefinedPushDownMethodsImpl* java file with some boiler-plate CCI code generated.

EJB Creates skeleton for the *UserDefinedPushDownMethodsImpl* java file with some boiler-plate EJB code generated.

JAX-RPC

Creates skeleton for the *UserDefinedPushDownMethodsImpl* java file with some boiler-plate JAX-RPC code generated.

WSIF Creates skeleton for the *UserDefinedPushDownMethodsImpl* java file with some boiler-plate WSIF code generated.

Custom

Creates empty skeleton for the *UserDefinedPushDownMethodsImpl* java file.

- c. Set additional attributes required for each *backEndType*. You are prompted for attributes according to which *backEndType* you specified.

CCI Set your bean's *connectionSpecClassName* attribute to the appropriate fully-qualified J2EE Connector Architecture (JCA) *ConnectionSpec* class name. For example, the class name for use with CICS ECI is *com.ibm.connector2.cics.ECIConnectionSpec*.

Set the *interactionSpecClassName* attribute to the appropriate fully-qualified J2EE Connector Architecture (JCA) *InteractionSpec* class name. For example, the class name for use with CICS ECI is *com.ibm.connector2.cics.ECIInteractionSpec*.

EJB Set the *ejbRefName* value that specifies the *java:comp* name for a reference to another EJB. For example, *ejb/AccountBackend*.

JAX-RPC

Set the *wSDLServiceFileName* attribute to the fully-qualified path (inside a JAR file) to a WSDL file that defines the web service being used.

WSIF Set the *wSDLServiceFileName* attribute to the fully-qualified path (inside a JAR file) to a WSDL file that defines the web service being used.

Custom

No special attributes.

All backend types

All backend types have an *ejbToRAAdapterClassName* entry, which enables you to specify a fully-qualified class name of a custom *EJBToRAAdapter* implementation.

- d. If you defined any abstract methods for pushdown, define a *PushDownMethodElement* for each of them, and set the *backEndMethodName* attribute to the name of the corresponding method.

3. Deploy the EJB JAR.

- a. Run the code generation tool *CMPDeploy* against the assembled EJB JAR file. You can do this from within the IDE or from the command line. The command syntax is:

```
CMPDeploy <ejb.jar | application.ear> <-dir directory> <-overwrite> <-verbose>
```

ejb.jar | application.ear

The name of the JAR or EAR file you want to deploy.

-dir Optional. The name of the directory into which you want the output files placed.

If you do not specify *-dir* the tool outputs its files into a package-based directory structure rooted at the current directory.

-overwrite

Optional. Causes any existing *UserDefinedPushDownMethodsImpl* class to be overwritten by the one you are creating.

If you do not specify `-overwrite` the tool will abort the process with a warning message if the *UserDefinedPushDownMethodsImpl* class already exists.

-verbose

Optional. Causes tracing messages to display during processing.

If you do not specify the `-verbose` option the tool does not display tracing messages, but only displays a message at the end of the process specifying either what was emitted or else what problems were encountered.

This command accepts either an EAR or JAR file, and generates the *UserDefinedPushDownMethods* interface, and skeleton Java source code for the bean's *UserDefinedPushDownMethodsImpl* class. For example, if the EJB JAR contains a *procedural* CMP bean called *Account*, then **CMPDeploy Account.jar** emits a generic *AccountUserDefinedPushDownMethods* interface, and a backend specific *AccountUserDefinedPushDownMethodsImpl.java* class that implements the interface.

4. Edit the emitted implementation class. Enter the appropriate connectivity logic for your specified backend data store.
5. Compile the emitted files.
6. Zip/jar the compiled files into the EJB JAR file.
7. Run the EJB deployment tool against the EJB JAR file. For more information, see documentation on the EJB deployment tool in the WebSphere Business Integration Server Foundation InfoCenter..
8. Package the deployed EJB JAR file into an EAR file.
9. Install the application containing this EJB EAR file. During installation, when asked for the resource binding for the CMP bean, specify the JNDI name of a *J2CConnectionFactory* (instead of the usual case of specifying the JNDI name of a *JDBC DataSource*), In the case of a non-CCI-based CMP bean (one that does NOT use JCA connectors), you must install the sample relational adapter provided by WebSphere Business Integration Server Foundation Version 5.1, and specify a binding to it.

Example: Container Managed Persistence over anything -- CCI Usage Scenario

For a CMP bean whose fully-qualified name is *com.ibm.test.BeCashAcctBean* (whose bean implementation class is named *BeCashAcctBean*), the *BeCashAcctBeanUserDefinedPushDownMethods* interface would be in the **com.ibm.test.websphere_deploy** package, and the *BeCashAcctBeanUserDefinedPushDownMethodsImpl* for a *backEndType* of CCI would be in the **com.ibm.test.websphere_deploy.cci** package.

Example

```
/**
 * User-defined push-down method ejbCreate.
 *
 * @param bean Reference to the bean implementation class
 * @param connection The CCI connection to the back-end system. This
 * connection was previously obtained (by the EJBTORAAAdapter implementation)
 * via connectionFactory.getConnection(connectionSpec). For non-CCI based
 * beans, this will be null and the user will have to manually connect to
 * the back-end system.
 *
 * @exception javax.ejb.CreateException Thrown if an exception occurs creating bean data
 * @exception ResourceException Any other exceptions are wrapped in a
 * ResourceException.
 */
public void ejbCreate(BeCashAcctBean bean, Object connection)
    throws javax.ejb.CreateException, ResourceException {
```

```

if (tc.isEntryEnabled())
    Tr.entry(
        tc,
        "ejbCreate(BeCashAcctBean, Object)",
        new Object[] { bean, connection });

try {
    com.ibm.connector2.cics.ECIInteractionSpec iSpec =
        new com.ibm.connector2.cics.ECIInteractionSpec();
    iSpec.setFunctionName("BECASHAC");
    // Specify any additional iSpec settings here
    iSpec.setCommareaLength(1242);
    iSpec.setReplyLength(1242);
    iSpec.setInteractionVerb(
        com.ibm.connector2.cics.ECIInteractionSpec.SYNC_SEND_RECEIVE);

    WSSstreamableRecord input = helper.createCCISstreamableRecord();

    String inputBufferString = createInputBuffer("01", bean);
    input.setValue(inputBufferString);

    WSSstreamableRecord output =
        (WSSstreamableRecord) helper.executeCCIInteraction(
            connection,
            iSpec,
            input);

    String errorCode =
        processOutputBuffer("01", bean, output.toString());
    if (!errorCode.equals("00000000")) {
        if (errorCode.equals("00000014")) {
            if (tc.isDebugEnabled()) {
                Tr.debug(tc, "Record already exists");
            }
            javax.ejb.DuplicateKeyException d =
                new javax.ejb.DuplicateKeyException();
            throw d;
        } else {
            if (tc.isDebugEnabled()) {
                Tr.debug(
                    tc,
                    "Unknown error with return code " + errorCode);
            }
            javax.ejb.CreateException c =
                new javax.ejb.CreateException(
                    "Unknown error with return " + "code " + errorCode);
            throw c;
        }
    }
    if (tc.isDebugEnabled()) {
        Tr.debug(tc, "Record created");
    }
} catch (Exception e) {
    ResourceException re =
        helper.createResourceException(e, this.getClass());

    if (tc.isEntryEnabled())
        Tr.exit(tc, "ejbCreate(BeCashAcctBean, Object)", re);

    throw re;
}

//Note that if the return code from the back-end datastore interaction
//indicated that a user-defined exception (declared in the push-down
//method's signature) should be thrown, then throw it here.

```

```

if (tc.isEntryEnabled())
    Tr.exit(tc, "ejbCreate(BeCashAcctBean, Object)");
}
/**
 * This method will use the bean information to create an
 * input buffer string for WStreamableRecord input that
 * maps to the COBOL commarea.
 *
 * Commarea:
 * -----
 *
 * This program uses a commarea to pass information to the backend
 * program. The format and size of the commarea is as follows:
 *
 * Request type (2) - the method to invoke. Must have a leading
 * zero
 * Return Code1 (8) - CICS DFHRESP value or a string with the
 * word ERROR or ABEND
 * Return Code2 (8) - CICS DFHRESP2 value or a string with the
 * error type or abend code.
 * Trace flag (1) - If set to something other than space or hex
 * zero, trace information will be written to a
 * CICS TransientData queue.
 * Total Records (4) - Number of records found on a browse.
 * ResType (2) - The resType for the CashAccount record
 * AccountId (8) - The accountid for the CashAccount record
 * Balance (8) - The balance for the CashAccount record
 * Type (1) - The type for the CashAccount record
 * Utilities (1200) - The utility info for the CashAccount record
 *
 * @param request type of method to invoke
 * @param bean Reference to the bean implementation class
 *
 * @return string that contains the mapped commarea.
 */
public String createInputBuffer(String requestType, BeCashAcctBean bean) {
    if (tc.isEntryEnabled())
        Tr.entry(
            tc,
            "createInputBuffer(requestType, bean)",
            new Object[] { requestType, bean });
    /*
     * Initialize buffer
     */
    StringBuffer sb = new StringBuffer(1242);
    sb.append(requestType);
    // request type (01=create, 02=find, 03=update, 04=delete)
    sb.append("00000000"); // return code 1
    sb.append("00000000"); // return code 2
    sb.append(" "); // trace
    sb.append("0001"); // total records
    sb.append("99"); // resType
    sb.append("88888888"); // accountId
    sb.append("00000000"); // balance
    sb.append("1"); // type
    sb.append("A"); // utilities

    /*
     * Update values from bean
     */
    String resType = bean.getResType();
    int start = 25 - resType.length();
    int end = 25;
    sb.replace(start, end, resType); // resType

```

```

String accoutId = bean.getAccountId();
start = 33 - accoutId.length();
end = 33;
sb.replace(start, end, accoutId); // accountId

    /*
    * Update values for.ejbStore
    */
if (requestType.equals("03")) {
    String balance = String.valueOf(bean.getBalance());
    start = 41 - balance.length();
    end = 41;
    sb.replace(start, end, balance);
    String rType = bean.getUType();
    start = 41;
    end = 42;
    sb.replace(start, end, rType);
    String utility = bean.getUtility();
    start = 42;
    end = 42 + utility.length();
    sb.replace(start, end, utility);
}

if (tc.isDebugEnabled()) {
    Tr.debug(
        tc,
        "The created string buffer for execute is " + sb.toString());
}
if (tc.isEntryEnabled())
    Tr.exit(tc, "createInputBuffer(requestType, bean)");
return sb.toString();
}

/**
 * This method will use the bean information to parse the
 * output buffer string that maps to information returned
 * from the helper.executeCCIInteraction plus the
 * COBOL commarea.
 *
 *
 * Commarea:
 * -----
 *
 * This program uses a commarea to pass information to the backend
 * program. The format and size of the commarea is as follows:
 *
 * Request type (2) - the method to invoke. Must have a leading
 * zero
 * Return Code1 (8) - CICS DFHRESP value or a string with the
 * word ERROR or ABEND
 * Return Code2 (8) - CICS DFHRESP2 value or a string with the
 * error type or abend code.
 * Trace flag (1) - If set to something other than space or hex
 * zero, trace information will be written to a
 * CICS TransientData queue.
 * Total Records (4) - Number of records found on a browse.
 * ResType (2) - The resType for the CashAccount record
 * AccountId (8) - The accountid for the CashAccount record
 * Balance (8) - The balance for the CashAccount record
 * Type (1) - The type for the CashAccount record
 * Utilities (1200) - The utility info for the CashAccount record
 *
 * @param request type of method to invoke
 * @param bean Reference to the bean implementation class
 * @param output buffer that contains the commarea
 *
 * @return first error code.
 */

```

```

public String processOutputBuffer(
String requestType,
BeCashAcctBean bean,
String outputBuffer) {
if (tc.isEntryEnabled())
Tr.entry(
tc,
"processOutputBuffer(requestType, bean, outputBuffer)",
new Object[] { requestType, bean, outputBuffer });

String returnCode1 = outputBuffer.substring(25, 33);
String returnCode2 = outputBuffer.substring(33, 41);
if (requestType.equals("03") && returnCode1.equals("00000000")) {
String balance = outputBuffer.substring(56, 64);
bean.setBalance(Integer.parseInt(balance));
String utype = outputBuffer.substring(64, 65);
bean.setUType(utype);
String utility = outputBuffer.substring(65, outputBuffer.length());
bean.setUtility(utility);
if (tc.isDebugEnabled()) {
Tr.debug(tc, "Balance is " + balance);
Tr.debug(tc, "utype is " + utype);
Tr.debug(tc, "utility is " + utility);
}
}
if (tc.isEntryEnabled())
Tr.exit(
tc,
"processOutputBuffer(requestType, bean, outputBuffer)");

return returnCode1;
}

```

Example: Container Managed Persistence over anything -- EJB Usage Scenario

For a CMP bean whose fully-qualified name is `com.ibm.test.Account` (whose bean implementation class is named `AccountBean`), the `AccountBeanUserDefinedPushDownMethods` interface would be in the `com.ibm.test.websphere_deploy` package, and the `AccountBeanUserDefinedPushDownMethodsImpl` for a `backEndType` of EJB would be in the `com.ibm.test.websphere_deploy.ejb` package.

Example

```

public float computeAverageDailyBalance(Date arg1, Date arg2, AccountBean bean,
Object connection) throws AccountAccessException, ResourceException
{
AccountBackendLocalHome home = helper.getEJBHome("ejb/AccountBackend",
AccountBackendLocalHome.class);
//throws ResourceException on error

AccountBackend bean = null;
float returnValue = 0.0;
int id = bean.getId();
try {
bean = home.create();
returnValue = bean.computeAverageDailyBalance(id, arg1, arg2);
} catch (Exception e) {
throw helper.createResourceException(e, this.getClass());
}

if (returnValue == -1) throw new AccountAccessException();

return returnValue;
}

```

Managing EJB containers

Each application server can have a single EJB container; one is created automatically for you when the application server is created. The following steps are to be performed only as needed to improve performance after the EJB application has been deployed.

1. Adjust EJB container settings.
2. Adjust EJB cache settings.

If adjustments do not improve performance, consider adjusting access intent policies for entity beans, reassembling the module, and redeploying the module in the application.

EJB container settings

Use this page to configure and manage a specific EJB container.

To view this administrative console page, click **Servers > Application Servers > *serverName* > EJB Container**.

Passivation directory

Specifies the directory into which the container saves the persistent state of passivated stateful session beans.

Beans are passivated when the number of active bean instances becomes greater than the cache size specified in the container configuration. When a stateful bean is passivated, the container serializes the bean instance to a file in the passivation directory and discards the instance from the bean cache. If, at a later time, a request arrives for the passivated bean instance, the container retrieves it from the passivation directory, deserializes it, returns it to the cache, and dispatches the request to it. If any step fails (for example, if the bean instance is no longer in the passivation directory), the method invocation fails.

For a cluster of servers that span multiple systems in a sysplex and have stateful session beans with an activation policy of *Transaction* deployed in them, the passivation directory must reside on a hierarchical file system (HFS) that is shared across the multiple systems.

Inactive pool cleanup interval

Specifies the interval at which the container examines the pools of available bean instances to determine if some instances can be deleted to reduce memory usage.

Data type	Integer
Units	Milliseconds
Range	0 to 2 147 483 674

Default datasource JNDI name

Specifies the JNDI name of a data source to use if no data source is specified during application deployment. This setting is not applicable for EJB 2.x-compliant CMP beans.

Servlets and enterprise beans use *data sources* to obtain these connections. When configuring a container, you can specify a default data source for the container. This data source becomes the default data source used by any entity beans installed in the container that use container-managed persistence (CMP).

The default data source for a container is *secure*. When specifying it, you must provide a user ID and password for accessing the data source.

Specifying a default data source is optional if each CMP entity bean in the container has a data source specified in its configuration. If a default data source is not specified and a CMP entity bean is installed in the container without specifying a data source for that bean, applications cannot use that CMP entity bean.

Initial state

Specifies the execution state requested when the server first starts.

Data type	String
Default	Started
Range	Valid values are Started and Stopped

EJB container system properties

In addition to the settings accessible from the administrative console, you can set the following system property by command-line scripting:

com.ibm.websphere.ejbcontainer.poolSize

Specifies the size of the pool for the specified bean type. This property applies to stateless, message-driven and entity beans. If you do not specify a default value, the container defaults of 50 and 500 are used.

Set the pool size for a given entity bean as follows:

```
beantype=min,max[:beantype=min,max...]
```

beantype is the J2EE name of the bean, formed by concatenating the application name, the # character, the module name, the # character, and the name of the bean (that is, the string assigned to the <ejb-name> field in the bean's deployment descriptor). *min* and *max* are the minimum and maximum pool sizes, respectively, for that bean type. Do not specify the square brackets shown in the previous prototype; they denote optional additional bean types that you can specify after the first. Each bean-type specification is delimited by a colon (:).

Use an asterisk (*) as the value of *beantype* to indicate that all bean types are to use those values unless overridden by an exact bean-type specification somewhere else in the string, as follows:

```
*=30,100
```

To specify that a default value be used, omit either *min* or *max* but retain the comma (,) between the two values, as follows (split for publication):

```
SMAApp#PerfModule#TunerBean=54,  
:SMAApp#SModule#TypeBean=100,200
```

You can specify the bean types in any order within the string.

EJB cache settings

Use this page to configure and manage the cache for a specific EJB container. To determine the cache absolute limit, multiply the number of enterprise beans active in any given transaction by the total number of concurrent transactions expected. Then, add the number of active session bean instances. You can use the Tivoli Performance Viewer to view bean performance information.

To view this administrative console page, click **Servers > Application Servers > *serverName* > EJB Container > EJB Cache Settings**.

Cleanup interval

Specifies the interval at which the container attempts to remove unused items from the cache in order to reduce the total number of items to the value of the cache size.

The cache manager tries to maintain some unallocated entries that can be allocated quickly as needed. A background thread attempts to free some entries while maintaining some unallocated entries. If the thread

runs while the application server is idle, when the application server needs to allocate new cache entries, it does not pay the performance cost of removing entries from the cache. In general, increase this parameter as the cache size increases.

Data type	Integer
Units	Milliseconds
Range	0 to 2 147 483 674
Default	3000

Cache size

Specifies the number of buckets in the active instance list within the EJB container.

A bucket can contain more than one active enterprise bean instance, but performance is maximized if each bucket in the table has a minimum number of instances assigned to it. When the number of active instances within the container exceeds the number of buckets, that is, the cache size, the container periodically attempts to reduce the number of active instances in the table by passivating some of the active instances. For the best balance of performance and memory, set this value to the maximum number of active instances expected during a typical workload.

Data type	Integer
Units	Buckets in the hash table
Range	Greater than 0. The container selects the next largest prime number equal to or greater than the specified value.
Default	2053

Container interoperability

Container interoperability describes the ability of WebSphere Application Server clients and servers at different versions to successfully negotiate differences in native Enterprise JavaBeans (EJB) Version 1.1 finder methods support and Java 2 Platform, Enterprise Edition (J2EE) Version 1.3 compliance.

At one time, there were significant interoperability problems among WebSphere Application Server, versions 4.0.x and 3.5.x distributed, and Version 4.0.x for zSeries. The introduction of interoperable versions of some class types solved these problems for distributed versions 3.5.6, 4.0.3, and 5 as well as for zSeries Version 4.0.x.

Older 4.0.x and 3.5.x client and application server versions do not support the interoperability classes, which makes them uninteroperable with versions that use the classes. The system property *com.ibm.websphere.container.portable* remedies this situation by enabling newer versions of the application server to turn off the interoperability classes. This lets a more recent application server return class types that are interoperable with an older client.

Depending on the value of *com.ibm.websphere.container.portable*, application servers at versions 5 and later, 4.0.3 and later, and 3.5.6 and later, return different classes for the following:

- Enumerations and collections returned by EJB 1.1 finder methods
- EJBMetaData
- Handles to:
 - Entity beans
 - Session beans
 - Home interfaces

If the property is set to *false*, application servers return the old class types, to enable interoperability with versions 3.5.5 and earlier, and 4.0.2 and earlier. If the property is set to *true*, application servers return the new classes.

Instructions for setting the `com.ibm.websphere.container.portable` property are in the release notes for versions 3.5.6 and later, and 4.0.3 and later. The following tables show interoperability characteristics for various version combinations of application servers and clients as well as default property values for each combination.

Interoperability of Version 3.5.x client with Version 5 (and later) application server

Clients at Version 3.5.5 and earlier are not interoperable with Version 5 and later servers when using:

- EJBMetaData
- Enumerations returned by EJB 1.x finder methods
- Handles to entity beans

If you would like to use updated Handle classes in EJB 2.x-compliant beans but have one of the older clients (versions 3.5.5 and earlier) installed, set the system property `com.ibm.websphere.container.portable.finder` to `false`. With this setting in place, the Version 5 application server uses the updated handles but returns the enumerations and collections that were used in the earlier clients.

To interoperate with Version 5 application servers, you must upgrade all Version 3.5.x clients to Version 3.5.6 or later.

Interoperability of Version 5 (and later) client with Version 3.5.x application server

Client at Version 5 and later, using this function	Application server at Version 3.5.6, property true	Application server at Version 3.5.6, property false (default)	Application server at Version 3.5.5 and earlier
EJBMetaData	Does not work across domains	Works	Does not work
Handle to session bean	Works	Works	Does not work
Handle to entity bean	Does not work across domains	Does not work across domains	Does not work across domains
Enumeration returned by EJB 1.x finder method	Works	Works	Works

Interoperability of Version 4.0.x client with Version 5 (and later) application server

Ideally, all 4.0.x clients that use Version 5 application servers should be at Version 4.0.3 or later.

Version 5 and later application servers return the interoperability class types by default (`true`). This can cause interoperability problems for distributed clients at versions 4.0.1 or 4.0.2. In particular, problems can occur with collections and enumerations returned by EJB 1.1 finder methods.

Although it is strongly discouraged, you can set `com.ibm.websphere.container.portable` to `false` on a Version 5 and later application server. This causes the application server to return the old class types, providing interoperability with clients at Version 4.0.2 and earlier. This is discouraged because:

- The Version 5 application server instance would become non-J2EE 1.3 compliant with regard to handles, home interface handles, and EJBMetaData.
- EJB 1.x finder methods return collection and enumeration objects that do not originate from `ejbportable.jar`.
- Interoperability restrictions still exist with the property set to `false`.
- Version 5 and later client handles to entity beans and home interfaces do not work across domains for the server you set to `false`.

If you would like to use updated Handle classes in EJB 2.x-compliant beans but have one of the older clients (versions 4.0.2 and earlier) installed, set the system property

com.ibm.websphere.container.portable.finder to false. With this setting in place, the Version 5 and later application server uses the updated handles but returns the enumerations and collections that were used in the earlier clients.

Interoperability of client at Version 4.0.2 and earlier with Version 5 (and later) application server

Client at Version 4.0.2 and earlier, using this function	Application server at Version 5 and later, property true (default)	Application server at Version 5 and later, property false
EJBMetaData	Does not work	Works for 4.0.2 client
Handle to session bean	Does not work	Works
Handle to entity bean	Does not work	Does not work across cells
Enumeration returned by EJB 1.x finder method	Does not work	Works
Collection returned by EJB 1.x finder method	Does not work	Works
Handle to home interface	Does not work	Does not work across cells

If you would like to use updated Handle classes in EJB 2.x-compliant beans but have one of the older clients (versions 3.5.5 and earlier, and 4.0.2 and earlier) installed, set the system property com.ibm.websphere.container.portable.finder to false. With this setting in place, the Version 5 and later server uses the new Handle classes but returns the older enumeration and collection classes.

Interoperability of client at Version 4.0.3 and later with Version 5 and later application server

Clients at Version 4.0.3 and later work well with Version 5 and later application servers. However, if you set the com.ibm.websphere.container.portable to false, client handles to entity beans and home interfaces do not work across domains for the server you set to false.

Client at Version 4.0.3 and later, using this function	Application server at Version 5 and later, property true (default)	Application server at Version 5 and later, property false
EJBMetaData	Works	Works
Handle to session bean	Works	Works
Handle to entity bean	Works	Does not work across cells
Enumeration returned by EJB 1.x finder method	Works	Works
Collection returned by EJB 1.x finder method	Works	Works
Handle to home interface	Works	Does not work across cells

Interoperability of Version 5 and later client with Version 4.0.x application server

Clients at Version 5 and later work well with Version 4.0.3 application servers if you set com.ibm.websphere.container.portable to true. Client handles to entity beans and home interfaces do not work across domains for any Version 4.0.3 server with com.ibm.websphere.container.portable at the default value, false. Version 5 client handles to application servers at Version 4.0.2 and earlier also have restrictions.

Client at Version 5 and later, using this function	Application server at Version 4.0.3, property true	Application server at Version 4.0.3, property false (default)	Application server at Version 4.0.2 or earlier
EJBMetaData	Works	Works	Works for 4.0.2 server only

Client at Version 5 and later, using this function	Application server at Version 4.0.3, property true	Application server at Version 4.0.3, property false (default)	Application server at Version 4.0.2 or earlier
Handle to session bean	Works	Works	Works
Handle to entity bean	Works	Does not work across domains	Does not work across domains
Enumeration returned by EJB 1.x finder method	Works	Works	Works
Collection returned by EJB 1.x finder method	Works	Works	Works
Handle to home interface	Works	Does not work across domains	Does not work across domains

Interoperability of zSeries Version 4.0.x client with Version 5 and later application server

The only valid configuration for container interoperability with zSeries Version 4.0.x clients is the default configuration for the Version 5 application server.

Interoperability of Version 5 and later client with zSeries Version 4.0.x application server

Version 5 clients should work with a zSeries Version 4.0.x application server with the correct interoperability fixes described in the zSeries documentation. The interoperability characteristics should be the same as for a Version 4.0.3 distributed application server with the property set to true.

Client at Version 5 and later, using this function	zSeries application server at Version 4.0.x
EJBMetaData	Works
Handle to session bean	Works
Handle to entity bean	Works
Enumeration returned by EJB 1.x finder method	Works
Collection returned by EJB 1.x finder method	Works
Handle to home interface	Works

Interoperability of the handle formats in WebSphere Application Server, Version 5 and Version 5.0.1

Applications that attempt to persist handles to enterprise beans and **EJBHome** needed to subclass **ObjectInputStream** in WebSphere Application Server, Version 5. This action was required so that the subclass **ObjectInputStream** could utilize the context class loader to resolve the classes for enterprise beans and **EJBHome** stubs.

In addition, handles created and persisted in WebSphere Application Server, Version 5 only work with objects that have an unchanged remote interface. If the remote interface is changed, the handle is no longer valid because the stub is serialized inside the handle and its serial Version UID changes if the remote interface changes.

This release introduces a new handle persistence mechanism that avoids the implementation drawbacks of the previous version. However, if handles are used for this WebSphere Application Server deployment, you should consider the following issues when applying this update, future WebSphere Application Server Fix Packs and EJB Container cumulative fixes for WebSphere Application Server, Version 5.

If a WebSphere Application Server, Version 5 persisted handle or home handle is encountered by a WebSphere Application Server, Version 5.0.1 system, it can be read and utilized. In addition, it will be

converted to WebSphere Application Server, Version 5.0.1 format if it is re-persisted. The WebSphere Application Server, Version 5.0.1 format cannot be read by a WebSphere Application Server, Version 5 system unless PQ72184 is applied.

Problems arise when handles are persisted and shared across systems that are not at the WebSphere Application Server, Version 5.0.1 level or later. However, a Version 5 system can receive a handle from Version 5.0.1 remotely through a call to get a handle on an enterprise bean or a getHomeHandle on an **EJBHome**. The remote call will succeed, however, any attempt to persist it on the Version 5 system will have the same limitations regarding the use of ObjectInputStream and changes in remote interface invalidating the persisted handle.

When your application stores handles persistently and shares this persistence with multiple clients or application servers, apply WebSphere Application Server, Version 5.0.1 or PQ72184 to both the client and server systems at the same time. Failure to do so can result in the inability of these systems to read the handle data stored by upgraded systems. Also, handles stored by the WebSphere Application Server, Version 5 can force the applications of the updated system to still subclass ObjectInputStream. Applications using the WebSphere Application Server Enterprise, Version 5 scheduler and process choreographer, are affected by these changes. These users should update their Version 5 systems at the same time with either Version 5.0.1 or PQ72184.

If the applications store handles in the session context, or locally in a file on the same system, that is not shared by other applications, on different systems, they might be able to update their systems individually, rather than all at once. If Client Container and thin client applications do not share persisted handle data, they can be updated as needed as well. However, handles created and persisted in WebSphere Application Server, Version 5, Version 4.0.3 and later (with the property flag set), or Version 3.5.7 and later (with the property flag set) are not usable if either the home or the remote interface changes.

If any WebSphere Application Server, Version 3.5.7 or Version 4.0.3 and later enables the system property `com.ibm.websphere.container.portable` to **true**, any handles to objects on that server have the same interoperability limitations. In addition, if any WebSphere Application Server, Version 3.5.7 and later or Version 4.0.3 applications store a handle obtained from a WebSphere Application Server, Version 5 or Version 5.0.1, the same restrictions apply, regarding the need to subclass ObjectInputStream and the usability of handles after a change to the remote interface is made.

Replication of the Http Session and Handles

This note applies to you if you place Handles to Homes or EJBs, or EJB or EJBHome references in the Http Session in your application and you use Http Session Replication. If you intend to replicate a mixed environment of Version 5.0.0 and Version 5.0.1 or 5.0.2 machines you should first apply the latest Version 5.0.0 container cumulative e-fix to the Version 5.0.0 machines before allowing the Version 5.0.1 or 5.0.2 server into the typology. The reason for this is that Version 5.0.0 servers are not able to understand the persisted Handle format used on the Version 5.0.1 and 5.0.2 server. This is similar to the case of Version 5.0.0 and Version 5.0.1 or 5.0.2 systems trying to use a shared database, mentioned above. But in this case, it is the Http Session object and not the database providing the persistence.

Top Down Deployment Mapping

The size of the Handle objects has grown due to the fix put in to allow serialization and deserialization to occur without the previous requirements of subclassing the ObjectInputStream and so on. Top down deployment of an object that contains EJB and EJBHome references create a database table ddl that has a field of 1000 bytes of VARCHAR for BITDATA which will contain the Handle. It might be that your object's Handle does not fit in the 1000 byte default field, and you might need to adjust this to a higher value. You might try increments of 250 bytes, that is, 1250, 1500, and so on.

Deploying EJB modules

Assemble one or more EJB modules, assemble one or more Web modules, and assemble them into a J2EE application.

1. Prepare the deployment environment. (See "Preparing to host applications" in the Information Center.)
2. Deploy the application.
3. **5.1+** Update the configuration for each EJB module as needed for the deployment environment.
4. For information about the EJB deployment tool, see the WebSphere Business Integration Server Foundation InfoCenter.

The next step is to test and debug the module. (See "Diagnosing and fixing problems" in the Information Center.)

EJB module collection

Use this page to manage the EJB modules deployed in a specific application.

To view this administrative console page, click **Applications > Enterprise Applications > *applicationName* > EJB modules**. Click the check boxes to select one or more of the EJB modules in your collection.

URI

When resolved relative to the application URL, this specifies the location of the module's archive contents on a file system. The URI matches the <ejb> or <web> tag in the <module> tag of the application deployment descriptor.

EJB module settings

Use this page to configure and manage a specific deployed EJB module.

Note: You cannot start or stop an individual EJB module for modification. You must start or stop the appropriate application entirely.

To view this administrative console page, click **Applications > Enterprise Applications > *applicationName* > EJB modules > *moduleName***.

URI

When resolved relative to the application URL, this specifies the location of the module archive contents on a file system. The URI must match the URI of a ModuleRef URI in the deployment descriptor of the deployed application (EAR).

Alternate DD

Specifies a deployment descriptor to be used at run time instead of the one installed in the module.

Starting weight

Specifies the order in which modules are started when the server starts. The module with the lowest starting weight is started first.

Data type	Integer
Default	5000
Range	Greater than 0

Enterprise beans: Resources for learning

Use the following links to find relevant supplemental information about enterprise beans. The information resides on IBM and non-IBM Internet sites, whose sponsors control the technical accuracy of the information.

These links are provided for convenience. Often, the information is not specific to this product but is useful all or in part for understanding the product. When possible, links are provided to technical papers and Redbooks that supplement the broad coverage of the release documentation with in-depth examinations of particular product areas.

View links to additional information about:

- Planning, business scenarios, and IT architecture
- Programming model and decisions
- Programming instructions and examples
- Programming specifications

Planning, business scenarios, and IT architecture

- Mastering Enterprise JavaBeans

A comprehensive treatment of Enterprise JavaBeans (EJB) programming in nonprintable form (PDF). One must be registered to download the PDF, but registration is free. Information about purchasing a hardcopy is available on the Web site.

- *Enterprise JavaBeans* by Richard Monson-Haefel (O'Reilly and Associates, Inc.: Third Edition, 2001)

Programming model and decisions

- Read all about EJB 2.0

A comprehensive overview of the specification.

- The J2EE Tutorial

This set of articles by Sun Microsystems covers several EJB-related topics, including the basic programming models, persistence, and EJB Query Language.

Programming instructions and examples

- Rules and Patterns for Session Facades

EJB programming practice: Fronting entity beans with a session-bean facade.

- WebSphere Application Server Development Best Practices for Performance and Scalability

Programming practice for enterprise beans and other types of J2EE components.

- Optimistic Locking in IBM WebSphere Application Server 4.0.2

Examples of the effect of optimistic concurrency on application behavior. Although the paper is based on a previous version of this product, the data access issues discussed in it are current.

This paper does not seem to be available directly by URL. To view this paper, visit the specified URL and search on "optimistic locking"

Programming specifications

- What's new in the Enterprise JavaBeans 2.0 Specification?

You can also download the specification itself from this URL.

- JavaTM 2 Platform: Compatibility with Previous Releases

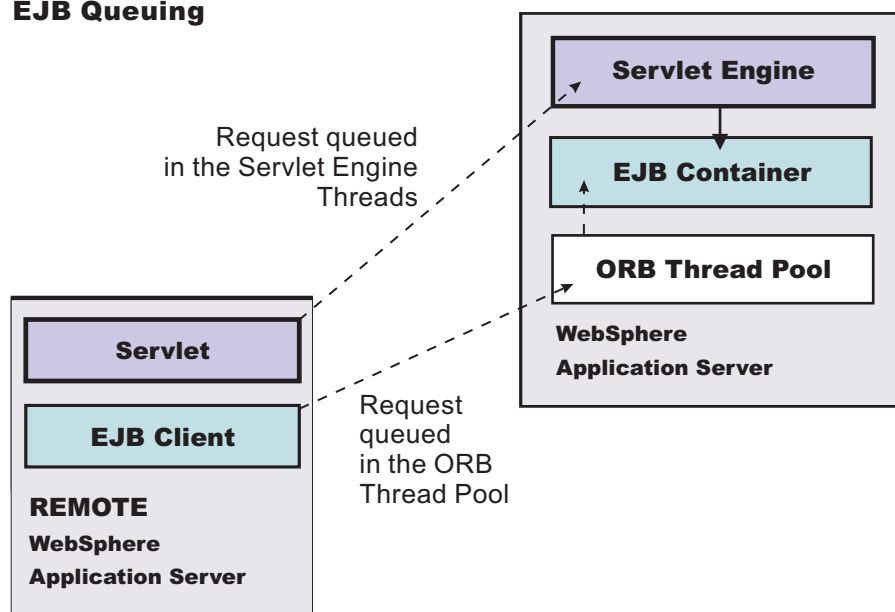
This Sun Microsystems article includes both source and binary compatibility issues.

EJB method Invocation Queuing

Method invocations to enterprise beans are only queued for remote clients, making the method call. An example of a remote client is an enterprise Java bean (EJB) client running in a separate Java virtual machine (JVM) (another address space) from the enterprise bean. In contrast, no queuing occurs if the EJB client, either a servlet or another enterprise bean, is installed in the same JVM on which the EJB method runs and on the same thread of execution as the EJB client.

Remote enterprise beans communicate by using the Remote Method Invocation over an Internet Inter-Orb Protocol (RMI-IIOP). Method invocations initiated over RMI-IIOP are processed by a server-side object request broker (ORB). The thread pool acts as a queue for incoming requests. However, if a remote method request is issued and there are no more available threads in the thread pool, a new thread is created. After the method request completes the thread is destroyed. Therefore, when the ORB is used to process remote method requests, the EJB container is an open queue, due to the use of unbounded threads. The following illustration depicts the two queuing options of enterprise beans.

EJB Queuing




The following are two tips for queueing enterprise beans:

- **Analyze the calling patterns of the EJB client.**

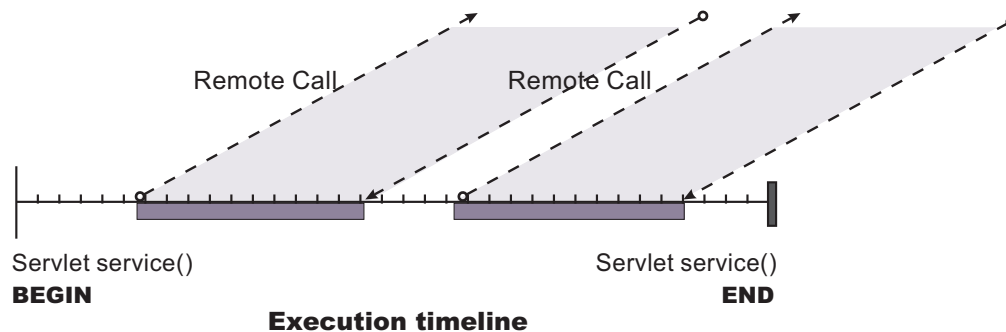
When configuring the thread pool, it is important to understand the calling patterns of the EJB client. If a servlet is making a small number of calls to remote enterprise beans and each method call is relatively quick, consider setting the number of threads in the ORB thread pool to a value lower than the Web

container thread pool size value.

Short-lived EJB calls



Longer-lived EJB calls



The degree to which the ORB thread pool value needs increasing is a function of the number of simultaneous servlets, that is, clients, calling enterprise beans and the duration of each method call. If the method calls are longer or the applications spend a lot of time in the ORB, consider making the ORB thread pool size equal to the Web container size. If the servlet makes only short-lived or quick calls to the ORB, servlets can potentially reuse the same ORB thread. In this case, the ORB thread pool can be small, perhaps even one-half of the thread pool size setting of the Web container.

- **Monitor the percentage of configured threads in use.**

Tivoli Performance Viewer shows a metric called *percent maxed*, which is used to determine how often the configured threads are used. A value that is consistently in the double-digits, indicates a possible bottleneck at the ORB. Increase the number of threads.

Chapter 5. Using extended messaging in applications

These topics provide information about implementing WebSphere enterprise applications that use extended messaging.

WebSphere Application Server supports asynchronous messaging as a method of communication based on the Java Message Service (JMS) programming interface. Extended messaging extends the base JMS support, support for EJB 2.0 message-driven beans, and the Enterprise Java Bean (EJB) component model, to use the existing container-managed persistence and transactional behavior.

Using extended messaging, you can build enterprise beans that can provide messaging services along with methods that implement business logic. The enterprise beans can use the standard JMS styles of messaging (point-to-point and publish/subscribe). However, with extended messaging, the JMS usage is simplified, because JMS support is managed by the extended messaging service. This helps to effectively separate business logic from the messaging infrastructure. The use of data mapping enables messages to drive existing or new enterprise beans as though they were invoked from any EJB client.

You can use WebSphere Studio Application Developer Integration Edition to develop applications that use extended messaging. You can use the WebSphere Application Server runtime tools, like the administrative console, to deploy and administer applications that use extended messaging.

For more information about implementing WebSphere enterprise applications that use extended messaging, see the following topics:

- Extended messaging - overview
- Extended messaging - application usage scenarios
- Extended messaging - components
- Designing an enterprise application to use a message-driven bean
- Developing an enterprise application to use a message-driven bean
- Deploying an enterprise application to use a message-driven bean
- Configuring extended messaging service resources
- Troubleshooting extended messaging

Extended messaging - overview

Extended messaging extends the base JMS support, support for EJB 2.0 message-driven beans, and the Enterprise Java Bean (EJB) component model, to use the existing container-managed persistence and transactional behavior.

In addition to providing such *container-managed messaging*, extended messaging provides new types of enterprise beans and administrative objects for messaging, and new functionality like data mapping and late response handling. (The abbreviation, CMM, for the term *container-managed messaging* is sometimes used to represent extended messaging.)

Extended messaging uses the bean-managed messaging implementation to provide the JMS interfaces, which ensures that both bean-managed and extended messaging use consistent JMS support.

An application that uses extended messaging can receive messages by using a receiver bean, either by the `onMessage()` method of a message-driven bean or by a stateless session bean that polls for a message from a named destination. With extended messaging and a message-driven bean, code within the bean can use the message to invoke business logic, as either a method within the same bean or another enterprise bean. Both the incoming message and the invocation of the receiver bean can be included within the scope of a transaction. For outbound messages, an application calls a sender bean that turns a method call into a JMS message that is then sent asynchronously. These message beans are implemented as enterprise beans by WebSphere Application Server. Application developers can create

these message beans by using WebSphere Studio Application Developer Integration Edition, although they can be created without the help of WebSphere Studio.

With extended messaging, the JMS usage is simplified, because JMS support is managed by the extended messaging service. This helps to effectively separate business logic from the messaging infrastructure. Also, the use of data mapping enables messages to drive existing or new enterprise beans as though they are invoked from any EJB client. WebSphere Studio enables the types of message beans that support extended messaging to be developed easily and hides the messaging infrastructure from developers.

For more conceptual information about extended messaging, see the following topics:

- Extended messaging - application usage scenarios
- Extended messaging - components
- Receiving messages with extended messaging
- Sending messages with extended messaging
- Data mapping with extended messaging
- Handling late responses with extended messaging
- Transactional support with extended messaging
- Exception handling with extended messaging

Extended messaging - receiving messages

To receive messages, applications that use extended messaging use a receiver bean, which can be a message-driven bean or a session bean:

- A receiver bean (deployed as a message-driven bean) is invoked when a message arrives at a JMS destination for which a listener is active.
- An application-callable receiver bean (deployed as a session bean) polls a JMS destination until a message arrives, gets the parsed message as an object, and can use getter methods to retrieve the message data.




Figure 1. Receiving messages with extended messaging. This figure shows an application calling a receiver bean (as a session bean) to receive messages from the JMS destination defined on an input port. The application also calls the `ReplySender()` method of the receiver bean to send a reply to the original message received. For more information about what is shown in this figure, see the text that accompanies this figure.

When a receiver bean gets a message, it can invoke another method passing either the JMS Message, or a set of parameters extracted from the message content. The invoked method can be contained in the receiver bean or in another enterprise bean (which is the preferred application structure). If data mapping is used, the method invoked by a receiver bean is unaware of the original JMS message.

In addition to receiving messages, extended messaging enables applications to reply to received messages in either of the following ways:

- **Sending a synchronous reply.** In this mode, which can be used by only receiver beans deployed as message-driven beans, the reply from the method invoked by the receiver bean is mapped to a JMS message and sent as a reply to the original message, using the replyTo field in the JMS header as the target destination.
- **Sending an asynchronous reply.** In this mode, which cannot be used receiver beans deployed as message-driven beans, the application calls the ReplySender() method to send the reply message. If the reply is passed as a set of parameters to the ReplySender() method, the reply is mapped to a JMS message before being sent.

If a receiver bean gets a JMS message, then depending upon the programming model (associated with the receiver bean by WebSphere Studio), one of the following interactions occurs:

- **Receive a request and send no reply.**

The receiver bean provides an anonymous invocation in the form of a method call. The data passed to the method is either the JMS message (if no data mapping is invoked) or a set of parameters mapped from the JMS message. The receiver bean cannot return a reply to this invocation. This mode of interaction can be used with point-to-point or publish/subscribe messaging.

- **Receive a request and send a synchronous reply.**

If the receiver bean gets a message, it invokes another method either in the same bean or another enterprise bean. When the method returns, the data returned from that call is mapped to a JMS message and sent to the reply destination specified in the original request message. The type of reply destination (queue or topic) must be the same as the type used by the original request received.

- **Receive a request and send an asynchronous reply.**

For a receiver bean deployed as a message-driven bean, the reply is returned (using the bean's ReplySender method) to the replyTo destination specified in the original request message.

For a receiver bean deployed as a session bean, the reply is returned to the destination defined in the input port for the receiver bean.

In addition to the asynchronous model of this interaction, this mode of interaction enables a method to send multiple replies to a single invocation.

Extended messaging - sending messages

To send messages, applications that use extended messaging call a method on a sender bean. A sender bean turns its method invocation into a JMS message, then passes that message to JMS. If needed, the sender bean can retrieve a response message, then translate that message into a result value and return it to the caller. If data mapping is used, the method that invokes a sender bean is unaware of the original JMS message. The sender bean methods can use data mapping to build JMS messages from data passed on the method call.




Figure 2. Sending messages with extended messaging. This figure shows an application calling a sender bean to send messages to the JMS destination defined on an output port. The application also calls the `receiveResponse()` method of the sender bean to receive a reply to the original message sent. For more information about what is shown in this figure, see the text that accompanies this figure.

A sender bean is an enterprise bean (stateless session bean) that can be built by WebSphere Studio Application Developer. A sender bean should not contain any application logic, to help preserve the separation between the messaging and business logic.

Each method defined on a class that implements a sender bean has one of the following modes of interaction (which is defined when the sender bean is built). The interaction extends the sender interface to address the issue of synchronizing anonymous invocations.

- **Send a request and receive no response.**

To send a JMS message, an application invokes the sender bean's method. The caller of the sender bean's method cannot receive a response to the message sent. This mode of interaction can be used with point-to-point or publish/subscribe messaging.

- **Send a request and receive a synchronous response.**

To send a JMS message and wait for a synchronous response, an application invokes the sender bean's method. The sender bean uses the message sender (an interface to JMS provided by extended messaging) to send the message and, when the response is received, to return the response message to the caller of the sender bean. This mode of interaction can be used with point-to-point messaging only.

- **Send a request and receive a deferred response.**

To send a JMS message and wait for a deferred response, an application invokes the sender bean's method. The sender bean uses the message sender to send the message, then returns to the caller without waiting for the response. The response is returned by a generated `receiveResponse()` method. This mode of interaction enables an application to receive more than one response message, as the application is responsible for retrieving the responses. This mode of interaction can be used with point-to-point messaging only.

Extended messaging - data mapping

A message bean can use data mapping to map between a JMS message and data as arguments:

- With data mapping, the target method of a receiver bean for an anonymous invocation receives the contents of an asynchronous message as arguments. The extended messaging service parses the JMS message and maps from the JMS message to the method arguments. Similarly, to send a message, an application invokes a method on a sender bean with appropriate arguments. The extended messaging service packs appropriate arguments into a JMS message then sends the asynchronous request.
- Without data mapping, the target method of a receiver bean for an anonymous invocation receives a JMS message; no data mapping is performed by extended messaging. Similarly, to send a message, an application invokes a method of a sender bean with a JMS message.

If a developer selects data mapping when creating a sender or receiver bean, extended messaging uses the parameter properties specified on the sender or receiver bean method signatures to perform the data mapping.

Extended messaging - handling late responses

If an application uses a sender bean to send a message, it can optionally retrieve a response to the message. The sender bean can either wait for the response or defer retrieval of the response. Sometimes a response is delayed within the messaging infrastructure, and therefore the application cannot receive the response. Extended messaging can retrieve such a response message (referred to as a late-response message) when it does arrive and pass it to a message-driven bean provided by the application to handle late responses. The message-driven bean used to handle the late response is a standard EJB 2.0 message-driven bean or a receiver bean deployed as a message-driven bean. The deployed message-driven bean can then perform its processing on the message.

Late responses should not be considered normal application behaviour.

For extended messaging to handle late responses for an application, the sender bean must be deployed with the **Handle late responses** option enabled.

Definition of a late response

A late response occurs when the application is no longer able to retrieve responses to messages that it has sent, as follows:

- **Send with deferred response.**

The application (enterprise bean) repeatedly tries to retrieve a response until it ends. When the application no longer wants to retry to get a response, it can register a request for extended messaging to handle the late response, by calling a `registerLateResponse()` method on the sender bean.

- **Send with synchronous response handling.**

When the sender bean sends a message, it waits for the response. The result of this is that either the sender bean retrieves the response message or a timeout error occurs. If the system raises a timeout error, the application can no longer retrieve a response to the message. At this time the extended messaging service registers the the message for a late response.

Handling responses

Extended messaging handles responses in the following stages:

1. Registering an interest in having a late response retrieved when it is available.

To request the system to handle late responses for a sender bean, you deploy the sender bean with the **Handle late responses** extension to the deployment Descriptor.

If selected, the **Handle late responses** option defines that extended messaging should pass the response, when it becomes available, to the message-driven bean provided by the application to handle late responses. When the sender bean is deployed a specialized listener port is associated with the bean. This listener port is known as a handle late response listener port.

If the option is not selected, then the system does not handle late responses, and it is the application's responsibility to handle any late responses.

2. Starting a JMS listener to retrieve the message when it is available, which then drives the message bean to handle the JMS message.

The listener port must be defined with the following properties:

- The same JMS destination as specified as the JMS response destination on the output port used by the sender bean.
- A listener port extension with Handle Late Responses enabled.

You cannot use a temporary destination for late responses.

3. If a request is made to handle a late response, the extended messaging service immediately registers a `LateResponse` message request with the extended message consumer for the given listener port.

The message request is registered independently of any transaction context that the sender bean has. A request record (containing the MessageID of the late response) is added to the AsyncMessageLog log. When the message is eventually received, it is passed to the message-driven bean deployed against the specified late response ListenerPort.

Extended messaging - transactional support

The global transaction context is not flowed on asynchronous (messaging) requests, so the receipt of an asynchronous message cannot be part of some existing remotely-established transaction. Reliability in an asynchronous environment is built on the message provider's ability to guarantee a once-and-once-only message delivery.

Transactional support with extended messaging builds on, and extends, the transactional support provided with bean-managed messaging, as follows:

- Transactional support for receiving messages
- Transactional support for sending messages

Transactional support for receiving messages (receiver beans)

The extended messaging transactional behaviour for receiver beans depends upon whether the bean is a receiver bean or an application-callable receiver bean.

- For a receiver bean (deployed as a message-driven bean), incoming message receipts are defined by the Transaction attribute of the onMessage() bean method. Message-driven beans can use bean-managed transactions (BMT) or container-managed transactions (CMT). For message-driven beans using CMT there are only two supported transaction options: Not supported and Required. If a message is to be received within a transaction, the message-driven bean must be using CMT with the Transaction attribute set to Required.
- For a receiver bean as a session bean, the bean only supports container-managed transactions, and the behaviour is defined by the Transaction attribute of the receiver bean method.

Dealing with retries

Note: In the asynchronous environment of transaction processing, rolling back a message receipt means that the message is not removed from the source destination. Although this behaviour is desirable and correct, it causes the message retained on the source queue to be reprocessed until the transaction commits. For receiver beans, you can control this behaviour as follows:

- **Receiver bean.**

To limit the number of times that a transaction is retried, you can either rely on the facilities of the JMS provider or use the retry limit facility of the Message Listener:

- WebSphere MQ JMS support provides the ability to move the message to a backout queue and uses two queue attributes, the backout threshold and the backout-requeue queue, to perform this.
- The Message Listener retry count can be used to stop the listener processing the queue if the threshold is reached. The listener behaviour can be disabled by setting the retry count value higher than the JMS provider threshold value.

- **Application-callable receiver bean.**

To limit the number of times that a transaction is retried depends on the facilities of the JMS provider to move the message to a backout queue.

Transactional support for sending messages (sender beans)

The transactional behaviour for sending messages is defined by the Transaction attribute on the send method within the sender bean.

If the `send()` method is part of a transaction, then the sending of an outgoing message occurs within any currently active transaction. This means that the message is not transmitted until the transaction is committed. If no transaction is active when the request to send the message occurs, then the message is transmitted immediately.

The transactional behaviour where the mode of interaction for a sender bean method specifies a response (that is, either *Send message and receive synchronous response* or *Send message and receive deferred response*) depends upon the type of response, as follows:

Transactional behaviour for a synchronous response

The sending of the request message and the receipt of the response message cannot be performed inside a transaction, because they are both performed within the same method. Therefore, the `send` is always non-transactional, regardless of the transactional setting for the method. The `receive` is either transactional or not, depending upon the Transaction setting of the method.

Transactional behaviour for a deferred response

The response message is received by a `receiveResponse()` method, which can have a different transactional behaviour to the sender method inside the sender bean. The transaction containing the `send` command must commit before the response can be received. The transactional behaviour is specified on the `send` and `receive` methods of the Sender bean.

Extended messaging - exception handling

Extended messaging provides the following exception handling for receiver beans and sender beans:

- Error handling for receiver beans
- Error handling for sender beans

Error handling for receiver beans

The following error conditions can lead to extended messaging exceptions

- Formatting error parsing the message, when performing data mapping
- Exception thrown by the application method
- `CMMException` when sending the reply

Errors are always logged. If the application does not catch the exception, the default behavior is to roll back any active transaction. If the received message is rolled back, then it can be processed again. This can occur a number of times until the message causing the error is removed from the queue by the JMS provider. (For more information, see *Dealing with retries* in Transactional support with extended messaging.)

With data mapping, if a receiver bean is deployed as a message-driven bean and a `replyTo` destination is configured, then error messages are sent as replies to that destination.

Application enterprise beans that call receiver beans deployed as session beans need to handle the `CMMException` exception. `CMMException` is an application exception which is declared in the `throws` clause of the methods in the generated receiver bean.

Error handling for sender beans

The following error conditions can lead to extended messaging exceptions

- Constructing the JMS message when data mapping from the parameters to the message
- Creating a message sender and sending the message
- Getting the response and parsing the message content

Errors are always logged. If the application does not catch the exception, the default behavior is to roll back any active transaction.

Application enterprise beans that call sender beans need to handle the `CMMException` exception, which is declared in the `throws` clause of the methods in the generated sender bean.

Extended messaging - application usage scenarios

Applications can use extended messaging to receive and send messages in a variety of ways:

- To receive messages, applications that use extended messaging use a receiver bean (deployed as a message-driven bean) or an application-callable receiver bean (deployed as a session bean):
 - A receiver bean is invoked when a message arrives at a destination for which a listener is active.
 - An application-callable receiver bean polls a destination (defined by an input port) until a message arrives or a timeout occurs.

In addition to receiving messages, extended messaging enables applications to send replies in response to the received messages.

For more conceptual information about receiving messages, see [Receiving messages with extended messaging](#).

- To send messages, applications that use extended messaging call sender bean methods. The sender bean sends messages to the target destination defined by an output port. The sender bean methods can be passed either a JMS message or a number of parameters that are mapped by extended messaging into a JMS message. Whether or not data mapping is used is specified when the application is developed.

In addition to sending messages, applications can choose to receive a response to the message, and can handle any responses either synchronously or asynchronously. If a response is not received in time, then the system can handle the late response by directing the message to a message-driven bean.

For more conceptual information about sending messages, see [Sending messages with extended messaging](#).

- Applications can combine receiving and sending messages in a variety of different ways. For example, a receiver bean deployed as a message-driven bean can forward the message by calling a sender bean. The receiver bean can give message data to sender bean in either of the following ways:
 - The receiver bean can pass the JMS message to the sender bean, which forwards that message.
 - The receiver bean can extract data from the initial message and send that data to the sender bean. The sender bean can then map the data values to a new JMS message, which it forwards.

The application can receive a response to the message that it sent, and then can send the message received or a new message built from data in the message received, as a response to the original message.

Also, data mapping can be used to hide the JMS message structure from the application logic. For more information about data mapping, see [Extended messaging - data mapping](#)

Extended messaging - components

Extended messaging builds on the base support for JMS messaging and message-driven beans provided by WebSphere Application Server. The new messaging components for extended messaging are referred to as the Message Bean package.

Components for receiving messages

The following components, shown in the figure [Components for receiving messages](#), are used to receive messages:

Receiver bean

An application that uses extended messaging can receive messages by using a receiver bean (using the `onMessage()` method of a message-driven bean) or an application-callable receiver bean (a stateless session bean that polls for a message from a named destination). Both receiver beans and application-callable receiver beans can receive and process asynchronous messages, and optionally return selected data as a response message.

Input port

An input port specifies the properties needed by receiver beans as session beans, by defining the following information:

- Information about the source destination for the message to be received
- Information about how to select and handle the message received
- Optional information about a reply destination, which is used if a reply is expected and replyTo information is not present in the JMSHeader of the message received.

A receiver bean as a deployed message-driven bean uses the associated listener port, so does not need an input port. For more information about message-driven beans and listener ports, see Message-driven beans - components.

For more conceptual information about receiving messages, see Receiving messages with extended messaging.




Figure 3. Components for receiving messages. This figure shows an application calling a receiver bean (as a session bean) to receive messages from the JMS destination defined on an input port. The application also calls the ReplySender() method of the receiver bean to send a reply to the original message received. For more information about what is shown in this figure, see the text that accompanies this figure.

Components for sending messages

The following components, shown in the figure Components for sending messages, are used to send messages:

Sender bean

Bean instances of a sender bean (also known as a message sender bean) can send asynchronous messages. The sender bean methods can be passed either a JMS message or a number of parameters that are mapped by extended messaging into a JMS message, which is then passed to JMS.

Output port

An output port specifies the properties needed by sender beans, to define the destination for the message being sent, and other optional properties if a response is expected. It is associated with the Sender Bean at deployment time and contains the following information:

- Information about the target destination for the message to be sent
- Information about how to select and handle the message to be sent
- Information about the destination used for the response.

For more conceptual information about sending messages, see [Sending messages with extended messaging](#).

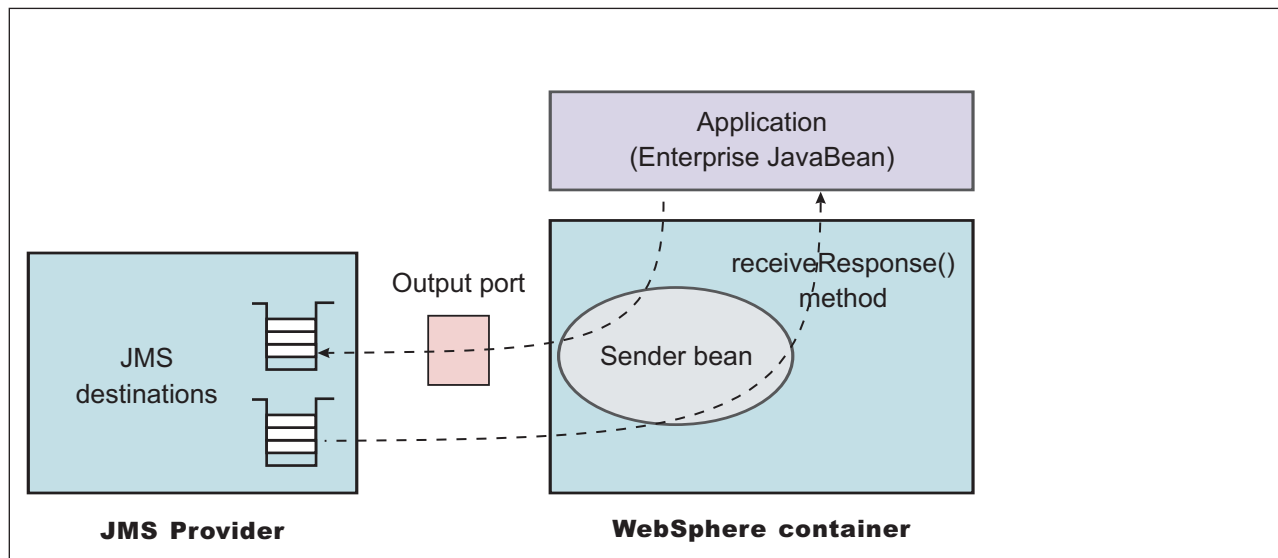


Figure 4. Components for sending messages. This figure shows an application calling a sender bean to send messages to the JMS destination defined on an output port. The application also calls the `receiveResponse()` method of the sender bean to receive a reply to the original message sent. For more information about what is shown in this figure, see the text that accompanies this figure.

Designing an enterprise application to use extended messaging

This topic describes things to consider when designing an enterprise application to use extended messaging.

The design of JMS-usage for applications that use extended messaging is the same as the design for JMS and message-driven beans, except that the JMS-usage is simplified because JMS support is managed by the extended messaging service. For design considerations related to JMS and message-driven beans, see the following topics:

-
-

The extra design consideration for applications that use extended messaging are as follows. For more detail, see the related topics.

1. For a receiver bean, decide whether to use a message-driven bean or stateless session bean.

Message-driven bean

You can use a deployed message-driven bean as a receiver bean, to automatically handle messages received at the associated listener port. As with any message-driven bean, when a message is received on the JMS destination monitored by the listener port, the message is passed to the `onMessage()` method of the message-driven bean.

You need to develop and deploy the message-driven bean, and configure its associated listener port, separately from the extended messaging tasks.

Stateless session bean

You can use a stateless session bean as a receiver bean, to poll for messages on a named destination associated with an input port.

You need to develop and deploy the session bean separately from the extended messaging tasks, but configure the associated input port as part of the extended messaging tasks.

2. Decide whether or not you want to use data mapping. If you call the methods of sender and receiver beans with data arguments, you need to use data mapping to construct the JMS messages needed. For data mapping, you need to decide what data arguments need to be specified as properties on the sender or receiver bean method signatures.

For a receiver bean deployed as a message-driven bean, you can define the mapping behavior if a data exception is caught by extended messaging. That is, you define whether a message should be flowed back if a ReplyTo destination is defined in the JMS message header.

3. Decide whether or not you want to handle late responses.

A sender bean can optionally retrieve a response to messages sent. If a response is delayed within the messaging infrastructure, the bean cannot receive the response. Extended messaging can retrieve such a response message (referred to as a late-response message) when it does arrive and pass it to a message-driven bean provided by the application to handle late responses. To handle late responses, you need to develop and deploy a standard EJB 2.0 message-driven bean that contains a registerLateResponse() method, and associate it with a listener port to be used to receive late responses.

Developing an enterprise application to use extended messaging

This topic describes how to develop an enterprise application to use extended messaging.

This task description assumes that developers are using the WebSphere Studio Application Developer to develop the application code (receiver and sender beans).

To develop an enterprise application to use extended messaging, complete the following steps:

1. Creating the Enterprise Application project. Because the sender and receiver beans used for extended messaging are EJB 2.0 enterprise beans, you must first have created a J2EE 1.3 Enterprise Application project for which extended messaging beans will be created.
 - a. Ensure that you have selected 1.3 as the highest J2EE version that is to be used in WebSphere Studio. For example: **Window-> Preferences... J2EE preferences-> Select the highest J2EE version that is to be used-> 1.3**
 - b. Create a J2EE 1.3 Enterprise Application project, as described in the WebSphere Studio article "Creating an Enterprise Application project."

2. Creating the application code.

To create the application code, use WebSphere Studio to generate the sender and receiver beans needed by the application, by completing one or more of the following subtasks as described in the WebSphere Studio Extended Messaging documentation:

- "Creating a sender bean"
- "Creating a receiver bean"
- "Creating an application-callable receiver bean"
- "Creating a sender bean and receiver bean"
- "Creating a sender bean and application-callable receiver bean"

The result of this stage is an enterprise bean, containing code automatically generated for extended messaging, that can be assembled into an .EAR file for deployment.

3. Assembling and packaging the application for deployment.

You can use WebSphere Studio to assemble and package the application for deployment.

The following aspects are specific to extended messaging:

- a. Configure a message selector for a receiver bean.
- b. Associate the JNDI names for sender and receiver beans with output and input ports.
- c. Specify the timeout for a sender bean response.
- d. Configure that a sender bean is to handle late responses and identify the listener port to be used for late responses.

The result of this task is an .EAR file, containing an application enterprise bean with code for extended messaging, that can be deployed in WebSphere Application Server.

For information about deploying an application to use extended messaging, see [Deploying an enterprise application to use extended messaging](#)

Deploying an enterprise application to use extended messaging

This topic describes how to deploy an enterprise application to use extended messaging.

This task description assumes that you have an .EAR file, which contains an application enterprise bean with code for extended messaging, that can be deployed in WebSphere Application Server.

The Application Install task is also a standard WebSphere Application Server task. As part of the install procedure you need to associate the Input and Output ports defined in System Management with the installed .EAR.

To deploy an enterprise application to use extended messaging, complete the following steps:

1. Use the administrative console to define and configure the extended messaging resources to be used by the application, as described in [Configuring extended messaging service resources](#). You should define the input ports for receiver beans, the output ports for sender beans, and listener port extensions for any sender beans that are to handle late responses.
2. Ensure that the deployment descriptor attributes for the sender and receiver beans match those of the extended messaging resources that you configured using the administrative console. The deployment descriptor values can be set when you generate the deployment code for the application, using WebSphere Studio Application Developer.
 - **5.1+** You can change the deployment descriptor values by using the Assembly Toolkit, as described in the following topics:
 - [Configuring deployment attributes for a receiver bean](#)
 - [Configuring deployment attributes for a sender bean](#).
3. If a sender bean is to handle late responses, deploy the message-driven bean to be used for late responses. For more information about deploying message-driven beans, see [Deploying an enterprise application to use message-driven beans](#).
4. Install the application into WebSphere Application Server.

This stage is a standard WebSphere Application Server task, as described in [Installing a new application](#).

When you install the application, you are prompted to specify the name of the listener port that the application is to use for late responses. Select the listener port, then click **OK**.

Configuring deployment attributes for a receiver bean

Use this task to configure the deployment attributes for a receiver bean for use with the extended messaging service.

You can specify these deployment attributes on each EJB method, as part of the deployment of the receiver bean. Changes to the deployment attributes override the values defined when the receiver bean was developed and deployment code was generated for the application.

You can configure the deployment attributes of an application by using the Deployment Descriptor Editor of WebSphere Studio Application Developer or the Assembly Toolkit.

This task description assumes that you have an EAR file for an application component, that can be deployed in WebSphere Application Server. For more details about using the Assembly Toolkit, see [Assembling applications with the Assembly Toolkit](#).

To configure the deployment attributes for a receiver bean, you can use the Assembly Toolkit to complete the following steps:

1. Start the Assembly Toolkit.
2. Create or edit the application EAR file. For example, to change attributes of an existing application, use the import wizard to import the EAR file into the Assembly Toolkit. To start the import wizard:
 - a. Click **File-> Import-> EAR file**
 - b. Click **Next**, then select the EAR file.
 - c. Click **Finish**
3. In the J2EE Hierarchy view, right-click the component instance, then click **Open With > Deployment Descriptor Editor**. A property dialog notebook for the component is displayed in the property pane.
4. In the property pane, click the Extended Messaging tab.
5. Specify appropriate deployment attributes: For example:
 - Method policies**
Add or remove method policies. To add a new method policy, complete the **Method policy name** and **Method policy description** fields, then click **Add**. To remove a method policy, select its name in the list displayed, then click **Remove**.
 - Input port**
For an application-callable receiver bean, this is the name of the input port to be used to receive messages.
 - Message selector**
For an application-callable receiver bean, this is a string used to select messages to be received.
6. Save your changes to the deployment descriptor.
 - a. Close the deployment descriptor editor.
 - b. When prompted, click **Yes** to indicate that you want to save changes to the deployment descriptor.
7. Verify the archive files.
8. Generate code for deployment for EJB modules or for enterprise applications that use EJB modules.
9. Optional: Test your completed module on a WebSphere Application Server installation. Right-click a module, click **Run on Server**, and follow the instructions in the displayed wizard. Note that **Run on Server** works on the Windows, Linux/Intel, and AIX operating systems only; you cannot deploy remotely from the Assembly Toolkit to a WebSphere Application Server installation on a UNIX operating system such as Solaris.

Important

Important: Use **Run On Server** for unit testing only. The Assembly Toolkit controls the WebSphere Application Server installation and, when an application is published remotely, the Toolkit overwrites the server configuration file for that server. Do not use on production servers. For instructions on remote testing, see the article “Setting Up a Remote WebSphere Application Server in WebSphere Studio V5” at http://www7b.boulder.ibm.com/wsdd/techjournal/0303_yuen/yuen.html.

After assembling your application, use a systems management tool to deploy the EAR file onto the application server that is to run the application; for example, using the administrative console as described in Deploying and managing applications.

Configuring deployment attributes for a sender bean

Use this task to configure the deployment attributes for a sender bean.

You can specify deployment attributes on each EJB method.

Changes to the deployment attributes override the values defined when the sender bean was developed and deployment code was generated for the application.

You can configure the deployment attributes of an application by using the Deployment Descriptor Editor of WebSphere Studio Application Developer or the Assembly Toolkit.

This task description assumes that you have an EAR file for an application component, that can be deployed in WebSphere Application Server. For more details about using the Assembly Toolkit, see *Assembling applications with the Assembly Toolkit*.

To change the deployment attributes for a sender bean, you can use the Assembly Toolkit to complete the following steps:

1. Start the Assembly Toolkit.
2. Create or edit the application EAR file. For example, to change attributes of an existing application, use the import wizard to import the EAR file into the Assembly Toolkit. To start the import wizard:
 - a. Click **File-> Import-> EAR file**
 - b. Click **Next**, then select the EAR file.
 - c. Click **Finish**
3. In the J2EE Hierarchy view, right-click the component instance, then click **Open With > Deployment Descriptor Editor**. For example, expand **EJB Modules-> ejb_module_instance-> Session Beans** then select the bean instance. A property dialog notebook for the component is displayed in the property pane.
4. In the property pane, click the Extended Messaging tab.
5. In the property pane, specify appropriate deployment attributes:
 - Output port**
This is the name of the output port to be used to send messages.
 - Enable handle late responses**
Select this checkbox if the sender bean is to handle late responses. If you select this checkbox, also specify the following properties: **ReplyTimeout** and **Late response handler listener port name**.
 - Late response handler listener port name**
For a sender bean that has been developed to handle late responses, this is the name of the listener port to be used for late responses.
 - Reply Timeout**
For a sender bean that has been developed to handle late responses, this is the time after which responses are considered late. This property is used if a response timeout is not specified on a sender method call.
6. Save your changes to the deployment descriptor.
 - a. Close the deployment descriptor editor.
 - b. When prompted, click **Yes** to indicate that you want to save changes to the deployment descriptor.
7. Verify the archive files.
8. Generate code for deployment for EJB modules or for enterprise applications that use EJB modules.
9. Optional: Test your completed module on a WebSphere Application Server installation. Right-click a module, click **Run on Server**, and follow the instructions in the displayed wizard. Note that **Run on Server** works on the Windows, Linux/Intel, and AIX operating systems only; you cannot deploy remotely from the Assembly Toolkit to a WebSphere Application Server installation on a UNIX operating system such as Solaris.

Important

Important: Use **Run On Server** for unit testing only. Assembly Server Toolkit controls the WebSphere Application Server installation and, when an application is published remotely, the Toolkit overwrites the server configuration file for that server. Do not use on production servers.

For instructions on remote testing, see the article "Setting Up a Remote WebSphere Application Server in WebSphere Studio V5" at http://www7b.boulder.ibm.com/wsdd/techjournal/0303_yuen/yuen.html.

Configuring extended messaging service resources

Use these tasks with the WebSphere Administrative console to configure resources needed by the extended messaging service and applications that use extended messaging.

You can use WebSphere Application Server system management to configure resources needed by the extended messaging service and applications that use extended messaging.

For more information about the tasks involved, see the following topics:

- Adding a new input port
- Adding a new output port
- Configuring an input port
- Configuring an output port
- Configuring a listener port to handle late responses

Adding a new input port

Use this task to add a new input port to WebSphere Application Server.

An input port is for use by an application that uses extended messaging.

During this task you configure the initial properties of the input port. You can later change the properties of the port, as described in [Configuring an input port](#).

To add a new input port, complete the following steps:

1. Start the WebSphere Administrative console.
2. In the navigation pane, select **Resources-> Extended messaging provider** This displays resources for extended messaging in the content pane.
3. In the Additional Properties table of the content pane, select **Input ports** This displays a list of the input ports in the content pane.
4. Click **New**.
5. Specify appropriate properties of the input port.
6. Click **OK**.
7. To save your configuration, click **Save** on the task bar of the Administrative console window.
8. To have the changed configuration take effect, stop then restart the application server.

Adding a new output port

Use this task to add a new output port to WebSphere Application Server, and configure its properties, for use by an application that uses extended messaging.

During this task you configure the initial properties of the output port. You can later change the properties of the port, as described in [Configuring an output port](#).

To add a new output port, complete the following steps:

1. Start the WebSphere Administrative console.
2. In the navigation pane, select **Resources-> Extended messaging provider** This displays resources for extended messaging in the content pane.
3. In the Additional Properties table of the content pane, select **Output ports** This displays a list of the output ports in the content pane.
4. Click **New**.
5. Specify appropriate properties of the output port.
6. Click **OK**.

7. To save your configuration, click **Save** on the task bar of the Administrative console window.
8. To have the changed configuration take effect, stop then restart the application server.

Configuring an input port

Use this task to change the properties of an input port for use by an application that uses extended messaging.

To change the properties of an input port, complete the following steps:

1. Start the WebSphere Administrative console.
2. In the navigation pane, select **Resources-> Extended messaging provider** This displays resources for extended messaging in the content pane.
3. In the Additional Properties table of the content pane, select **Input ports** This displays a list of the input ports in the content pane.
4. Select the input port that you want to change.
5. Specify appropriate properties of the input port.
6. Click **OK**.
7. To save your configuration, click **Save** on the task bar of the Administrative console window.
8. To have the changed configuration take effect, stop then restart the application server.

Configuring an output port

Use this task to change the properties of an output port for use by an application that uses extended messaging.

To change the properties of an output port, complete the following steps:

1. Start the WebSphere Administrative console.
2. In the navigation pane, select **Resources-> Extended messaging provider** This displays resources for extended messaging in the content pane.
3. In the Additional Properties table of the content pane, select **Output ports** This displays a list of the output ports in the content pane.
4. Select the output port that you want to change.
5. Specify appropriate properties of the output port.
6. Click **OK**.
7. To save your configuration, click **Save** on the task bar of the Administrative console window.
8. To have the changed configuration take effect, stop then restart the application server.

Extended messaging service settings

Use this page to enable or disable the extended messaging service.

The Extended Messaging Service provides run-time service for the support of extended messaging.

To view this administrative console page, click **Servers > Application Servers > *server_name* > Extended Messaging Service** .

Startup

Specifies whether the server will attempt to start the extended messaging service.

Default

Selected

Range**Selected**

When the application server starts, it attempts to start the extended messaging service automatically.

Cleared

The server does not try to start the extended messaging service. If extended messaging is to be used in applications that run on this server, the system administrator must start the extended messaging service manually or select this property then restart the server.

Late response handling extension collection

Use this page to view the configuration properties of late response handling extensions.

Late response handling extensions enable the handling of late responses with extended messaging

To view this administrative console page, click **Application Servers > server_name > Extended Messaging Service > Listener Port Extensions** .

Enabled:

Specifies whether the handling of late responses is enabled.

Range**Selected**

Handling of late responses is enabled.

Cleared

Handling of late responses is not enabled.

Request Interval:

Specifies the interval that elapses between checking for late responses.

Data type

Integer

Units

milliseconds

Default

5

Range

An integer number of milliseconds, greater than or equal to 0:

- 0 indicates that the late response handler continually checks for requests
- Other values are an integer number of milliseconds between checks for requests.

Request Timeout:

Specifies the duration of time after which to give up waiting for a response.

Data type

Integer

Units

seconds

Default

0

Range

An integer number of milliseconds, greater than or equal to -1:

- -1 indicates that requests to handle late responses are never discarded.
- Other values are an integer number of milliseconds after which requests are discarded.

Listener Ports:

Specifies the name of the listener port to be used to handle late responses.

Late response handling extension settings:

Use this page to configure late response handling extensions.

To view this administrative console page, click **Application Servers** > *server_name* > **Extended Messaging Service** > **Listener Port Extensions** > *extension_name* .

Enabled:

Specifies whether the handling of late responses is enabled.

Range

Selected

Handling of late responses is enabled.

Cleared

Handling of late responses is not enabled.

Request Interval:

Specifies the interval that elapses between checking for late responses.

Data type

Integer

Units

milliseconds

Default

5

Range

An integer number of milliseconds, greater than or equal to -1:

- -1 indicates that requests to handle late responses are never discarded.
- Other values are an integer number of milliseconds after which requests are discarded.

Request Timeout:

Specifies the duration of time after which to give up waiting for a response.

Data type

Integer

Units

seconds

Default

0

Range

An integer number of milliseconds, greater than or equal to -1:

- -1 indicates that requests to handle late responses are never discarded.
- Other values are an integer number of milliseconds after which requests are discarded.

Listener Ports:

Specifies the name of the listener port to be used to handle late responses.

Extended messaging provider settings

Use this page to manage extended messaging providers.

The extended messaging provider manages resources defined for use with extended messaging.

To view this administrative console page, click **Resources > Extended Messaging Providers** .

Name

The name of the resource provider.

Data type	String
Range	1 through 30 ASCII characters

Description

An optional description for the resource factory.

Data type	String
------------------	--------

Input port collection

Use this page to view the configuration properties of input ports..

An input port specifies the properties needed by receiver beans as session beans. Receiver beans as message-driven beans do not need an input port, because the properties needed are associated with the deployed message-driven bean and the Message Listener service.

To view this administrative console page, click **Resources > Extended Messaging Providers > Input Port** .

Name:

The name by which the input port is known for administrative purposes.

Data type	String
Units	En_US ASCII characters

JNDI Name:

The JNDI name for the resource.

Data type	String
------------------	--------

Description:

A description of the input port, for administrative purposes.

Data type	String
------------------	--------

Category:

A string that can be used to classify or group the resource.

Data type	String
Range	1 through 30 ASCII characters

JMS Connection Factory JNDI Name:

The JNDI name for the JMS connection factory to be used by the input port; for example, jms/connFactory1.

Data type String

JMS Destination JNDI Name:

The JNDI name for the destination to be used by the input port; for example, jms/destn1.

Data type String

JMS Acknowledgement Mode:

JMS acknowledgment mode to be used for acknowledging messages.

This property applies only to message-driven beans that use bean-managed transaction demarcation (**Transaction type** is set to Bean).

Default Range

Auto Acknowledge

Auto Acknowledge

The session automatically acknowledges a message in either of the following cases:

- When the session has successfully returned from a call to receive a message.
- When the session has called a message listener to process the message and received a successful response from that listener.

Dups OK Acknowledge

The session acknowledges only the delivery of messages. This is likely to result in the delivery of some duplicate messages if JMS fails, so it should be used only by consumers that are tolerant of duplicate messages.

Destination Type:

The type of the JMS resource.

Default Range

Queue

Queue The receiver bean receives messages from a queue destination.

Topic The receiver bean receives messages from a topic destination.

Subscription durability:

[Topic destinations only.] Specifies whether a JMS topic subscription is durable or non-durable.

Default Durable

Range**Durable**

A subscriber registers a durable subscription with a unique identity that is retained by JMS. Subsequent subscriber objects with the same identity resume the subscription in the state it was left in by the earlier subscriber. If there is no active subscriber for a durable subscription, JMS retains the subscription's messages until they are received by the subscription or until they expire.

Nondurable

Nondurable subscriptions last for the lifetime of their subscriber object. This means that a client sees the messages published on a topic only while its subscriber is active. If the subscriber is not active, the client is missing messages published on its topic.

Reply JMS Connection Factory JNDI Name:

JNDI name of the JMS Connection Factory to be used for replies.

Data type String

Reply JMS Destination JNDI Name:

JNDI name of the JMS Destination to be used for replies.

Data type String

Input port settings:

Use this page to configure an input port.

To view this administrative console page, click **Resources > Extended Messaging Providers > Input Port > *inputport_name*** .

Name:

The name by which the input port is known for administrative purposes.

Data type String
Units En_US ASCII characters

JNDI Name:

The JNDI name for the resource.

Data type String

Description:

A description of the input port, for administrative purposes.

Data type String

Category:

A string that can be used to classify or group the resource.

Data type String
Range 1 through 30 ASCII characters

JMS Connection Factory JNDI Name:

The JNDI name for the JMS connection factory to be used by the input port; for example, jms/connFactory1.

Data type String

JMS Destination JNDI Name:

The JNDI name for the destination to be used by the input port; for example, jms/destn1.

Data type String

JMS Acknowledgement Mode:

JMS acknowledgment mode to be used for acknowledging messages.

This property applies only to message-driven beans that use bean-managed transaction demarcation (**Transaction type** is set to Bean).

Default Auto Acknowledge
Range **Auto Acknowledge**
The session automatically acknowledges a message in either of the following cases:

- When the session has successfully returned from a call to receive a message.
- When the session has called a message listener to process the message and received a successful response from that listener.

Dups OK Acknowledge
The session acknowledges only the delivery of messages. This is likely to result in the delivery of some duplicate messages if JMS fails, so it should be used only by consumers that are tolerant of duplicate messages.

Destination Type:

The type of the JMS resource.

Default Queue
Range **Queue** The receiver bean receives messages from a queue destination.
Topic The receiver bean receives messages from a topic destination.

Subscription durability:

[Topic destinations only.] Specifies whether a JMS topic subscription is durable or non-durable.

**Default
Range**

Durable
Durable

A subscriber registers a durable subscription with a unique identity that is retained by JMS. Subsequent subscriber objects with the same identity resume the subscription in the state it was left in by the earlier subscriber. If there is no active subscriber for a durable subscription, JMS retains the subscription's messages until they are received by the subscription or until they expire.

Nondurable

Nondurable subscriptions last for the lifetime of their subscriber object. This means that a client sees the messages published on a topic only while its subscriber is active. If the subscriber is not active, the client is missing messages published on its topic.

Reply JMS Connection Factory JNDI Name:

JNDI name of the JMS Connection Factory to be used for replies.

Data type String

Reply JMS Destination JNDI Name:

JNDI name of the JMS Destination to be used for replies.

Data type String

Output port collection

Use this page to view the configuration properties of output ports.

The Output port defines the parameters required by the extended messaging sender bean. These properties define the destination for the message being sent, together with optional details if a response is expected.

To view this administrative console page, click **Resources > Extended Messaging Providers > Output Port** .

Name:

The name by which the output port is known for administrative purposes.

Data type String

JNDI Name:

The JNDI name for the output port.

Data type String

Description:

A description of the output port, for administrative purposes.

Data type String

Category:

A string that can be used to classify or group the resource.

Data type String

JMS Connection factory JNDI name:

The JNDI name for the JMS connection factory to be used by the output port; for example, jms/connFactory1.

Data type String
Units En_US ASCII characters
Range 1 through 30 ASCII characters

JMS Destination JNDI name:

The JNDI name for the destination to be used by the output port; for example, jms/destn1.

Data type String

JMS Delivery Mode:

Specifies whether all messages sent to the destination are persistent or non-persistent.

Default Persistent
Range **Persistent**
Messages put onto the destination are persistent.
Nonpersistent
Messages put onto the destination are not persistent.

JMS Priority:

The message priority for this queue destination.

Data type Integer
Default 4
Range 0 to 9.

JMS Time To Live:

The time in milliseconds after which messages on this queue expire.

Data type Integer
Units Milliseconds
Default 0

Range	0 to <i>n</i>
	0 messages never time out.
	<i>n</i> messages time out in <i>n</i> milliseconds.

JMS Disable Message I.D.:

Specifies that the system should not generate a JMS message ID.

Default	Cleared
Range	Selected The system does not generate message IDs.
	Cleared The system generates message IDs automatically.

JMS Disabled Message Time Stamp:

Specifies that the system should not generate a JMS message timestamp.

Default	Cleared
Range	Selected Message time stamps are added automatically to messages sent.
	Cleared Message time stamps are not added automatically to messages sent.

Response JMS Connection Factory JNDI name:

The JNDI name for the JMS connection factory to be used for response messages handled by the output port; for example, `jms/connFactory1`.

Data type	String
Units	En_US ASCII characters
Range	1 through 30 ASCII characters

Response JMS Destination JNDI name:

The JNDI name for the destination to be used for response messages handled by the output port; for example, `jms/destn1`.

Data type	String
------------------	--------

Output port settings:

Use this page to configure an output port.

An output port specifies the properties needed by sender beans to define the destination for the message being sent, and other optional properties if a response is expected. The output port is associated with the sender bean at deployment time.

To view this administrative console page, click **Resources > Extended Messaging Providers > Output Port > `outputport_name`** .

Name:

The name by which the output port is known for administrative purposes.

Data type String

JNDI Name:

The JNDI name for the output port.

Data type String

Description:

A description of the output port, for administrative purposes.

Data type String

Category:

A string that can be used to classify or group the resource.

Data type String

JMS Connection factory JNDI name:

The JNDI name for the JMS connection factory to be used by the output port; for example, jms/connFactory1.

Data type String
Units En_US ASCII characters
Range 1 through 30 ASCII characters

JMS Destination JNDI name:

The JNDI name for the destination to be used by the output port; for example, jms/destn1.

Data type String

JMS Delivery Mode:

Specifies whether all messages sent to the destination are persistent or non-persistent.

Default Persistent
Range **Persistent**
Messages put onto the destination are persistent.
Nonpersistent
Messages put onto the destination are not persistent.

JMS Priority:

The message priority for this queue destination.

Data type	Integer
Default	4
Range	0 to 9.

JMS Time To Live:

The time in milliseconds after which messages on this queue expire.

Data type	Integer
Units	Milliseconds
Default	0
Range	0 to <i>n</i> 0 messages never time out. <i>n</i> messages time out in <i>n</i> milliseconds.

JMS Disable Message I.D.:

Specifies that the system should not generate a JMS message ID.

Default	Cleared
Range	Selected The system does not generate message IDs. Cleared The system generates message IDs automatically.

JMS Disabled Message Time Stamp:

Specifies that the system should not generate a JMS message timestamp.

Default	Cleared
Range	Selected Message time stamps are added automatically to messages sent. Cleared Message time stamps are not added automatically to messages sent.

Response JMS Connection Factory JNDI name:

The JNDI name for the JMS connection factory to be used for response messages handled by the output port; for example, jms/connFactory1.

Data type	String
Units	En_US ASCII characters
Range	1 through 30 ASCII characters

Response JMS Destination JNDI name:

The JNDI name for the destination to be used for response messages handled by the output port; for example, jms/destn1.

Data type	String
------------------	--------

Troubleshooting extended messaging

Use this overview task to help resolve a problem that you think is related to the extended messaging service.

The extended messaging service uses the standard WebSphere Application Server RAS facilities. If you encounter a problem that you think might be related to the extended messaging service, complete the following stages:

1. Check for extended messaging service messages in the application server's SystemOut log at `was_home\logs\server_name\SystemOut`. Any error messages associated with the extended messaging service are labelled with EMSG. The error message indicates the nature of the problem and provides some detail. The extended messaging service issues EMSG error messages if it fails to initialize, parse its configuration file, or encounters some runtime error.
2. Check for more messages in the application server's SystemOut log. If the JMS server is running, but you have problems accessing JMS resources, check the SystemOut log file, which should contain more error messages and extra details about the problem.
3. Check the Release Notes for specific problems and workarounds. The section *Possible Problems and Suggested Fixes* of the Release Notes, available from the WebSphere Application Server library web site, is updated regularly to contain information about known defects and their workarounds. Check the latest version of the Release Notes for any information about your problem. If the Release Notes does not contain any information about your problem, you can also search the Technotes database on the WebSphere Application Server web site.
4. Check for problems with the WebSphere Messaging functions or message-driven beans. For more information about troubleshooting WebSphere Messaging, see the related topics listed at the bottom of this file.
5. Get a detailed exception dump for extended messaging. If the information obtained in the preceding steps is still inconclusive, you can enable the application server debug trace for the "Messaging" group to provide a detailed exception dump.

Extended Messaging: Resources for learning

Use the following links to find relevant supplemental information about Extended Messaging. The information resides on IBM and non-IBM Internet sites, whose sponsors control the technical accuracy of the information.

These links are provided for convenience. Often, the information is not specific to the IBM WebSphere Application Server product, but is useful all or in part for understanding the product. When possible, links are provided to technical papers and Redbooks that supplement the broad coverage of the release documentation with in-depth examinations of particular product areas.

View links to additional information about:

- "Extended Messaging: Resources for learning"
- "Extended Messaging: Resources for learning"
- "Extended Messaging: Resources for learning"

Programming model and decisions

- Sun's Java Message Service (JMS) specification documentation

Programming specifications

- Java Message Service API, 1.0.2
- Enterprise JavaBeans Technology Downloads & Specifications

Other

- WebSphere Application Server Enterprise Version 5 Overview: Extended J2EE Development Accelerators

- Listing of PDF files to learn about WebSphere Application Server Version 5
- Listing of all IBM WebSphere Application Server Redbooks
- Listing of all IBM WebSphere Application Server Whitepapers
- WebSphere Application Server Enterprise Edition 4.0: A Programmer's Guide

Chapter 6. Using asynchronous messaging

Use these topics to use asynchronous messaging with WebSphere Application Server, to enable enterprise applications to use JMS resources, message-driven beans, and extended messaging.

WebSphere Application Server supports asynchronous messaging as a method of communication based on the Java Message Service (JMS) programming interface.

The base JMS support enables WebSphere enterprise applications to exchange messages asynchronously with other JMS clients by using JMS destinations (queues or topics). An enterprise application can explicitly poll for messages on a destination.

WebSphere Application Server also provides a message listener service that enterprise applications can use to automatically retrieve messages from JMS destinations for processing by message-driven beans, without the application having to explicitly poll JMS destinations.

5.1 + With WBI Server Foundation, enterprise applications can use another level of functionality for asynchronous messaging, called extended messaging. The application server manages the messaging infrastructure, and extra standard types of messaging beans are provided to add functionality to that provided by message-driven beans. This level of functionality enables application developers to concentrate on the business logic to be implemented by the enterprise beans and to leave the messaging usage to standard messaging objects and configuration of the extended messaging service.

You can use the WebSphere administrative console to administer the WebSphere Application Server support for asynchronous messaging. For example, you can configure JMS providers and their resources, and can control the activity of the JMS server and the messaging services.

For more information about implementing WebSphere enterprise applications that use asynchronous messaging, see the following topics:

- An overview of WebSphere asynchronous messaging
- Implementing WebSphere enterprise applications that use JMS
- Implementing WebSphere enterprise applications that use message-driven beans
- Implementing WebSphere enterprise applications that use extended messaging

For more information about JMS, see the JMS documentation at <http://java.sun.com/products/jms/docs.html>.

Asynchronous messaging with WebSphere - an overview

WebSphere Application Server supports asynchronous messaging as a method of communication based on the Java Message Service (JMS) programming interface.

The base JMS support enables WebSphere J2EE applications to exchange messages asynchronously with other JMS clients by using JMS destinations (queues or topics). An J2EE application can explicitly poll for messages on a destination.

WebSphere Application Server also provides a message listener service that J2EE applications can use to automatically retrieve messages from JMS destinations for processing by message-driven beans, without the application having to explicitly poll JMS destinations.

5.1 + With WBI Server Foundation, J2EE applications can use another level of functionality for asynchronous messaging, called extended messaging. The application server manages the messaging infrastructure, and extra standard types of messaging beans are provided to add functionality to that provided by message-driven beans. This level of functionality enables application developers to

concentrate on the business logic to be implemented by the enterprise beans and to leave the messaging usage to standard messaging objects and configuration of the extended messaging service.

You can use the WebSphere administrative console to administer the WebSphere Application Server support for asynchronous messaging. For example, you can configure JMS providers and their resources, and can control the activity of the JMS server and the messaging services.

Styles of messaging in applications

Applications can use the following styles of asynchronous messaging:

Point-to-Point

Point-to-point applications use queues to pass messages between each other. The applications are called point-to-point, because a client sends a message to a specific queue and the message is picked up and processed by a server listening to that queue. It is common for a client to have all its messages delivered to one queue. Like any generic mailbox, a queue can contain a mixture of messages of different types.

Publish/subscribe

Publish/subscribe systems provide named collection points for messages, called topics. To send messages, applications publish messages to topics. To receive messages, applications subscribe to topics; when a message is published to a topic, it is automatically sent to all the applications that are subscribers of that topic. By using a topic as an intermediary, message publishers are kept independent of subscribers.

Both styles of messaging can be used in the same application.

Applications can use asynchronous messaging in the following ways:

One-way

An application sends a message, and does not want a response. This pattern of use is often referred to as a datagram.

Request / response

An application sends a request to another application and expects to receive a response in return.

One-way and forward

An application sends a request to another application, which sends a message to yet another application.

These messaging techniques can be combined to produce a variety of asynchronous messaging scenarios.

For more information about how such messaging scenarios are used by WebSphere enterprise applications, see the following topics:

- An overview of asynchronous messaging with JMS
- An overview of asynchronous messaging with message-driven beans
- An overview of extended messaging

For more information about these messaging techniques and the Java Messaging Service (JMS), see Sun's Java Message Service (JMS) specification documentation.

WebSphere Application Server cloning and WebSphere MQ clustering

This topic provides a summary of information about using WebSphere Application Server horizontal cloning with WebSphere MQ server clustering support. It describes a scenario that shows how the message listener service can be configured to take advantage of WebSphere MQ server clustering and provides some information about how to resolve potential runtime failures in the clustering scenario. The information in this topic is based on the scenario shown in the figure WebSphere Application Server horizontal cloning with WebSphere MQ clustered queues.

Note: WebSphere MQ server clustering is only available with the full WebSphere MQ product installed as the JMS provider.

For a WebSphere application server configured to use the extended messaging service, each JMS listener is used to retrieve messages from destinations defined to the server. In the following information the listener configurations are the same for each WebSphere application server. Each application server host contains a WebSphere application server and an WebSphere MQ server. If a host is only used to distribute messages, it only contains an WebSphere MQ server. There can be many servers defined in the configuration, although for simplicity the information in this topic is based on a scenario containing only three servers as shown in “WebSphere Application Server cloning and WebSphere MQ clustering” on page 160.

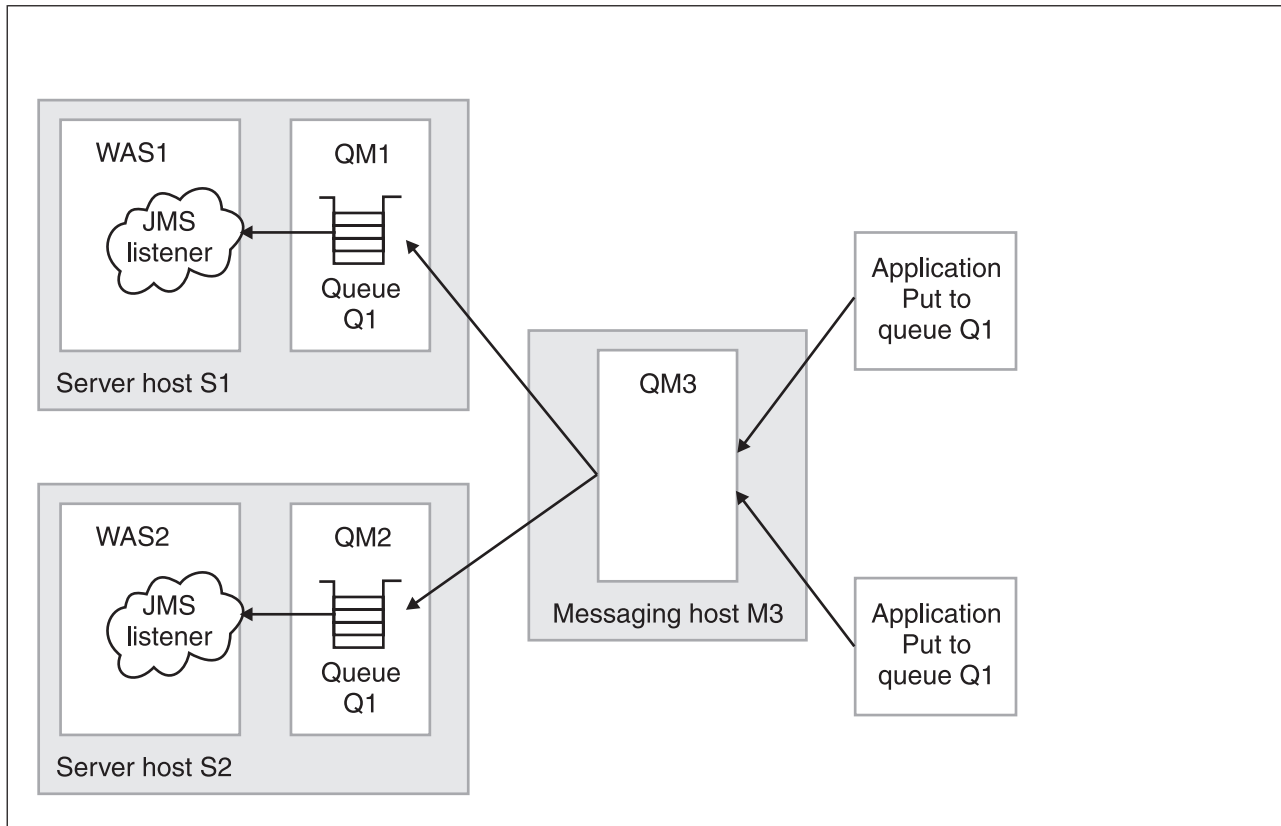


Figure 5. WebSphere Application Server horizontal cloning with WebSphere MQ clustered queues. This figure shows two WebSphere Application Server hosts, with horizontal clustering, and a messaging host used to distribute messages for WebSphere MQ server clustering. For more information, see the text that accompanies this figure.

The scenario shown in “WebSphere Application Server cloning and WebSphere MQ clustering” on page 160 comprises the following three hosts:

- Server host S1 contains the following servers:

WebSphere MQ server.

The server is defined to have a queue manager, QM1, and a local queue, Q1. The queue manager belongs to a cluster. The queue is populated by the WebSphere MQ server located on host M3. Applications can add messages directly to the queue, Q1, but would not be subjected to the control of the WebSphere MQ cluster.

WebSphere Application Server

This contains a cloned application server, WAS1, which is configured with a JMS listener. The listener is configured to retrieve messages from JMS destination Q1.

- Server host S2 contains the following servers:

WebSphere MQ server.

The server is defined to have a queue manager, QM2, and a local queue, Q1. The queue

manager belongs to the same cluster as QM1 on host S1. As with QM1, the queue is populated by the WebSphere MQ server located on host M3. Applications can add messages directly to the queue, Q1, but would not be subjected to the control of the MQ cluster.

WebSphere Application Server

This contains a cloned application server, WAS2 (identical to WAS1 on host S1), which is configured with a JMS listener. The listener is configured to retrieve messages from JMS destination Q1.

- Messaging host M3 contains the following servers:

WebSphere MQ server.

The server is defined to have a queue manager, QM3, which also belongs to the same cluster as QM1 and QM2. Applications add messages to the queue manager and queue Q1. The cluster to which this queue manager belongs causes messages to be distributed to all other queue managers in the cluster which have queue Q1 defined.

Note: Queue Q1 is not defined as a local queue on this host. If the queue was defined locally, then messages would remain on the server for local processing; messages would not be distributed by the queue manager cluster control to the other queue managers in the cluster that do have the queue defined.

This host does not have a WebSphere application server defined. All message retrieval processing is performed by the other two application servers on hosts S1 and S2.

Recovery scenarios

There are several failure scenarios that could occur with the clustering configuration; for example:

- WAS server failures.

In this scenario the failure of any single WebSphere application server results in the messages for the specified destination remaining on the queue, until the server is restarted.

- WebSphere MQ Queue Manager failures.

There are two different failures to consider:

1. Failure of a queue manager on the same host as a WebSphere application server (for example, failure of QM2 on host S2). In this case messages are delivered to the other available application servers, until the WebSphere MQ server is back online, when messages are processed as expected.
2. Failure of the messaging host M3 and its queue manager, QM3. In this case, the result of an outage is more significant because no messages are delivered to the other queue managers in the cluster. In a fully-deployed and scaled production system, host M3 would not be designed to be a single point of failure, and additional messaging servers would be added to the cluster configuration.

Using JMS and messaging in applications

Use these tasks to implement WebSphere J2EE applications that use JMS.

WebSphere Application Server supports asynchronous messaging as a method of communication based on the Java Message Service (JMS) programming interface.

The base JMS support enables WebSphere enterprise applications to exchange messages asynchronously with other JMS clients by using JMS destinations (queues or topics). An enterprise application can explicitly poll for messages on a destination.

Using the base support for JMS, you can build enterprise beans that use the JMS API directly to provide messaging services along with methods that implement business logic.

You can use the WebSphere administrative console to administer the JMS support of WebSphere Application Server. For example, you can configure JMS providers and their resources, and can control the activity of the JMS server.

For more information about implementing WebSphere enterprise applications that use JMS, see the following topics:

- An overview of WebSphere asynchronous messaging using JMS
- Administering WebSphere messaging
- Developing a J2EE application to use JMS
- Developing a JMS client
- Deploying a J2EE application to use JMS
- Tuning Java messaging service
- Troubleshooting WebSphere messaging

For more information about JMS, see the JMS documentation at <http://java.sun.com/products/jms/docs.html>.

An overview of WebSphere asynchronous messaging using JMS

WebSphere Application Server supports asynchronous messaging as a method of communication based on the Java Message Service (JMS) programming interface. JMS provides a common way for Java programs (clients and J2EE applications) to create, send, receive, and read asynchronous requests, as JMS messages.

This topic provides an overview of asynchronous messaging using JMS support provided by WebSphere Application Server. For more details about JMS, see Sun's Java Message Service (JMS) specification documentation.

The base support for asynchronous messaging using JMS, shown in the following figure, provides the common set of JMS interfaces and associated semantics that define how a JMS client can access the facilities of a JMS provider. This enables WebSphere J2EE applications, as JMS clients, to exchange messages asynchronously with other JMS clients by using JMS destinations (queues or topics). An J2EE application can use JMS queue destinations for point-to-point messaging and JMS topic destinations for publish/subscribe messaging. An J2EE application can explicitly poll for messages on a destination then retrieve messages for processing by business logic beans (enterprise beans).

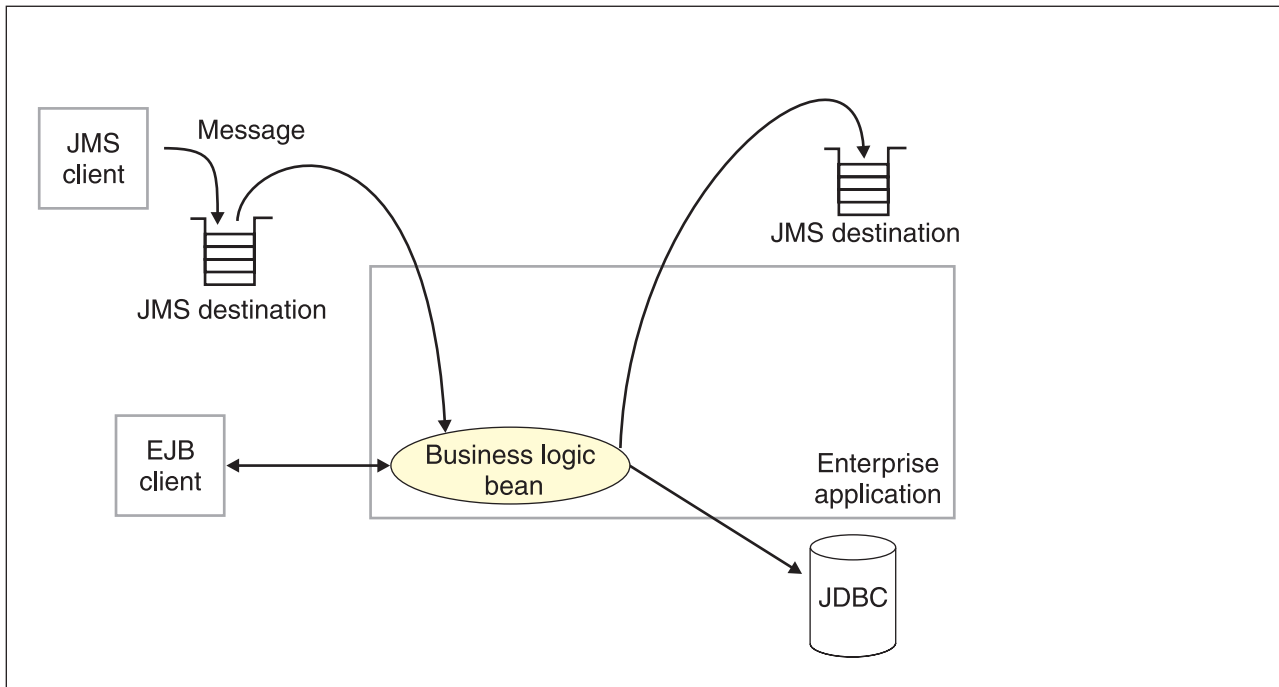


Figure 6. Asynchronous messaging using JMS. This figure shows an enterprise application polling a JMS destination to retrieve an incoming message, which it processes with a business logic bean. The business logic bean uses standard JMS calls to process the message; for example, to extract data or to send the message on to another JMS destination. For more information, see the text that accompanies this figure.

With the base JMS/XA support, the J2EE application uses standard JMS calls to process messages, including any responses or outbound messaging. Responses can be handled by an enterprise bean acting as a sender bean, or handled in the enterprise bean that receives the incoming messages. Optionally, this process can use two-phase commit within the scope of a transaction. This level of functionality for asynchronous messaging is called bean-managed messaging, and gives an enterprise bean complete control over the messaging infrastructure; for example, for connection and session pool management. The application server has no role in bean-managed messaging.

WebSphere Application Server also supports automatic asynchronous messaging using message-driven beans (a type of enterprise bean defined in the EJB 2.0 specification) and JMS listeners (part of the JMS application server facilities). Messages are automatically retrieved from JMS destinations, optionally within a transaction, then sent to the message-driven bean in an J2EE application, without the application having to explicitly poll JMS destinations. For more information about asynchronous messaging with message-driven beans, see An overview of asynchronous messaging with message-driven beans

5.1+ With WBI Server Foundation, J2EE applications can use another level of functionality for asynchronous messaging, called extended messaging. The application server manages the messaging infrastructure, and extra standard types of messaging beans are provided to add functionality to that provided by message-driven beans. This level of functionality enables application developers to concentrate on the business logic to be implemented by the enterprise beans and to leave the messaging usage to standard messaging objects and configuration of the extended messaging service.

WebSphere JMS support - components

The main components of WebSphere JMS support are shown in the following figure and described after the figure:

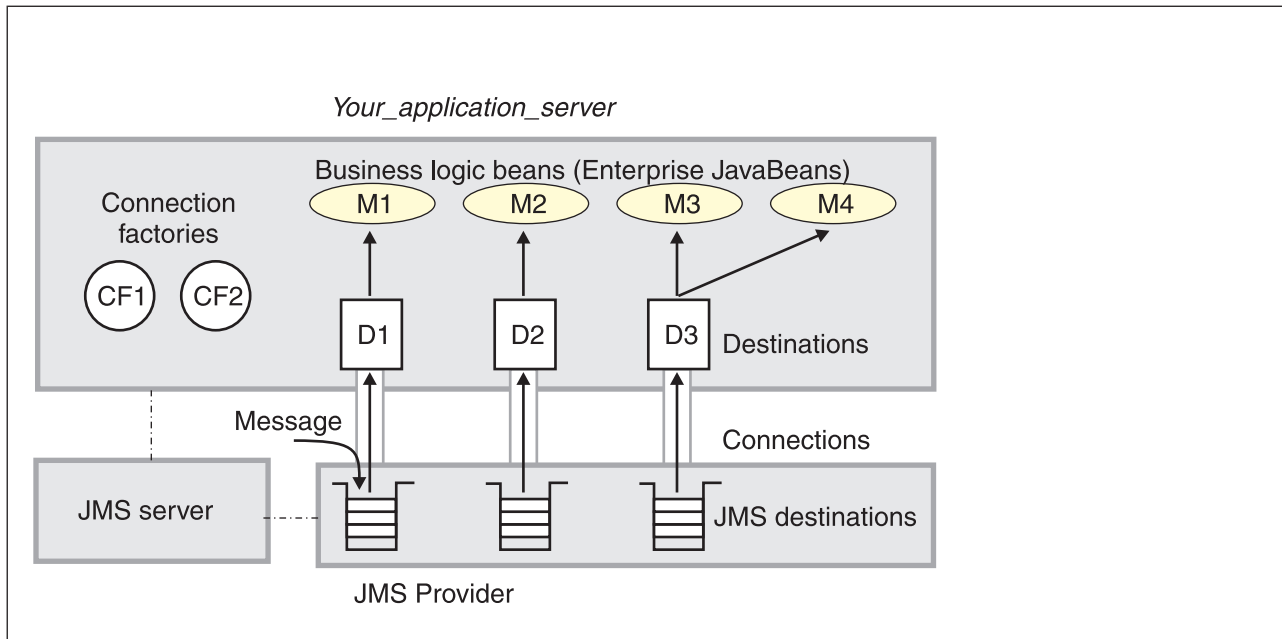


Figure 7. The main components of WebSphere JMS support. This figure shows the main components of WebSphere JMS support, from JMS provider through a connection to a destination, then to a WebSphere enterprise application (acting as a JMS client) that processes the message retrieved from the destination. For more information, see the text that accompanies this figure.

WebSphere Application Server supports asynchronous messaging based on the Java Messaging Service (JMS) of a JMS provider that conforms to the JMS specification version 1.0.2 and supports the Application Server Facility (ASF) function defined within that specification. WebSphere Application Server provides an embedded JMS provider and administration objects for WebSphere MQ as the JMS provider. You can use the embedded JMS provider, install WebSphere MQ JMS on top of the embedded WebSphere JMS, or install and configure another JMS provider.

JMS functions (of JMS providers) within the WebSphere Application Server administration domain are served by one or more JMS servers. There can be at most one JMS server on each node in the administration domain, and any application server within the domain can access JMS resources served by any JMS server on any node in the domain.

A connection factory is used to create connections with the JMS provider for a specific JMS queue or topic destination. Each connection factory encapsulates the configuration parameters needed to create a connection to a JMS destination.

A WebSphere J2EE application can explicitly poll for messages on a destination then retrieve messages for processing by business logic beans (enterprise beans).

The WebSphere Application Server support for message-driven beans and extended messaging builds on this base JMS support. For more information, see the related topics.

WebSphere MQ JMS connection pooling

To improve the overall performance of JMS within the system, the message listener service enables the connection pooling facility provided by the WebSphere MQ JMS implementation. This support does not affect the performance of a message listener, because it retains its connections while listening on a destination, but does affect the overall JMS system performance. When a connection is no longer required, WebSphere MQ can pool the connection then reuse it later instead of destroying it.

Note: This support is only available if WebSphere MQ is configured as the JMS provider.

To enable WebSphere MQ connection pooling and configure the characteristics of the WebSphere MQ connection pool, see [Enabling WebSphere MQ JMS connection pooling](#).

Administering WebSphere JMS support

Use these tasks with the WebSphere administrative console to manage JMS providers and their resources, and other runtime components of WebSphere JMS support.

You can use the WebSphere administrative console to configure the embedded WebSphere JMS provider or an WebSphere MQ JMS provider. If you install another JMS provider, you need to configure that JMS provider by using the tools and information provided with the JMS provider. For each JMS provider, you can configure the properties of JMS resources.

You can also use the WebSphere administrative console to configure and control other runtime components of WebSphere JMS support, including the following:

- The WebSphere JMS server
- The message listener service, listener ports, and the listener for each message-driven bean
- Input and output ports for extended messaging.

You can update the configuration data at any time, but if it is updated, the updates only take effect when the appropriate server is next started.

For information about the specific tasks used to administer WebSphere JMS support, see the following topics:

- [Installing and configuring a JMS provider](#)
- [Moving from the internal JMS provider to WebSphere MQ](#)
- [Securing messaging directories and log files](#)
- [Enabling security for the embedded WebSphere JMS provider](#)
- [Displaying administrative lists of JMS resources](#)
- [Managing JMS servers on Application Server](#)
- [Managing JMS servers in a deployment manager cell](#)
- [Configuring JMS provider resources](#)

For more information about JMS resources, see the JMS documentation at <http://java.sun.com/products/jms/docs.html>.

Installing and configuring a JMS provider

This topic describes the different ways that you can implement a JMS provider for use with WebSphere Application Server.

For IBM WebSphere Application Server to support bean-managed messaging, you need to install and configure one or more JMS providers that conform to the JMS specification version 1.0.2. To use message-driven beans the JMS provider must support the Application Server Facility (ASF) function defined within that specification.

You can install and use the Embedded Messaging Server option of WebSphere Application Server, install WebSphere MQ as the JMS provider, or install another “generic” JMS provider. If you install both embedded messaging and WebSphere MQ as JMS providers, for example, WebSphere applications can use JMS resources provided by both the embedded WebSphere JMS provider and the WebSphere MQ JMS provider.

WebSphere Application Server provides default JMS support with its Embedded Messaging Server component. This function is installed with WebSphere Application Server, administered through the administrative console and managed as part of the WebSphere Application Server runtime. This function is only accessible from WebSphere Application Server Web, EJB and client containers, and is not interoperable with WebSphere MQ. If access is required to heterogeneous non-JMS applications, WebSphere MQ clustering, or other WebSphere MQ functions, you should install WebSphere Application

Server without the Embedded Messaging Server component, and should install WebSphere MQ instead as the JMS provider. You are recommended to install and use the WebSphere Application Server Embedded Messaging Client for communication with the WebSphere MQ queue manager.

To provide the supported levels of messaging, you should apply the latest service for WebSphere Application Server and WebSphere MQ while a production application is in Quality-Assurance or Product-Validation Test.

For more information about WebSphere Application Server messaging scenarios, and the relationship between embedded messaging and WebSphere MQ, see the following articles:

- WebSphere Application Server solutions White paper: Selecting the most appropriate JMS provider for your applications (<ftp://ftp.software.ibm.com/software/websphere/resourcefinder/2318-00.pdf>). This white paper describes the different Java Message Service (JMS) options that are available from IBM for WebSphere Application Server applications. It compares the messaging technology built into WebSphere Application Server, Version 5 with IBM WebSphere MQ, IBM WebSphere Business Integration Event Broker and IBM WebSphere Business Integration Message Broker.
- JMS Topologies and Configurations with WebSphere Application Server and WebSphere Studio Version 5 (http://www7b.boulder.ibm.com/wsdd/library/techarticles/0310_barciabarci.html). This technical article helps you understand configuration options within WebSphere Application Server for Java Message Service (JMS) applications, including using embedded messaging and WebSphere MQ (formerly MQSeries). It covers setting up your development environment to support various JMS scenarios, development and runtime topologies, and scripting and security.

You can install and configure a JMS provider in one or more of the following ways:

- Installing WebSphere embedded messaging as the JMS provider.

Note:

- WebSphere embedded messaging as the JMS provider supports both queues (for point-to-point messaging) and topics (for publish/subscribe messaging).
 - You can install IBM WebSphere Application Server with embedded messaging on the same host as an existing WebSphere MQ installation, which must be at a supported level of MQ features.
 - You can install IBM WebSphere Application Server with embedded messaging then later install WebSphere MQ for use as a JMS provider.
- Installing WebSphere MQ as the JMS provider.

Note:

- You can install WebSphere MQ before IBM WebSphere Application Server. If you then want to install embedded messaging, you must ensure that the WebSphere MQ installation is at a supported level of MQ features.
- If you do not want to use the embedded WebSphere JMS provider, you can install IBM WebSphere Application Server without the **Embedded Messaging Server** option. You are recommended to install and use the WebSphere Application Server **Embedded Messaging Client**.
- You can install WebSphere MQ for use as a JMS provider on top of WebSphere Application Server embedded messaging; this results in a single JMS installation.
 - For point-to-point messaging WebSphere applications can continue to use WebSphere queue resources (through the embedded messaging JMS provider) or WebSphere MQ queue resources that you define to IBM WebSphere Application Server.
 - For publish/subscribe messaging, WebSphere applications can continue to use WebSphere topic resources (through the embedded messaging JMS provider) or WebSphere MQ topic resources that you define to IBM WebSphere Application Server, and which are provided by a Publish/Subscribe broker installed in addition to the base WebSphere MQ.

The preferred solution for publish/subscribe messaging with WebSphere MQ as the JMS provider is a full broker such as WebSphere MQ Event Broker. Failing this, you can install the MAOC SupportPac with WebSphere MQ.

- If you install WebSphere MQ as the JMS provider, you can use the WebSphere administrative console to administer the WebSphere MQ JMS provider resources, such as queue connection factories. However, you cannot administer MQ security, which is administered through WebSphere MQ.
- You can change from using WebSphere embedded messaging to WebSphere MQ as the JMS provider for WebSphere enterprise applications, as described in Moving from the embedded WebSphere JMS provider to WebSphere MQ.

For more information about scenarios and considerations for using WebSphere MQ with IBM WebSphere Application Server, see the White Papers and Red books provided by WebSphere MQ; for example, through the WebSphere MQ library Web page at <http://www-3.ibm.com/software/ts/mqseries/library/>

- Installing another JMS provider, which must conform to the JMS specification and, to use message-driven beans, support the ASF function. If you want to use a JMS provider other than the embedded WebSphere JMS provider or a WebSphere MQ JMS provider, you should complete the following steps:
 1. Installing and configuring the JMS provider and its resources by using the tools and information provided with the JMS provider.
 2. Defining the JMS provider to WebSphere Application Server as a generic JMS provider.

Note: You cannot use the WebSphere administrative console to administer the JMS provider or its resources.

To install a JMS provider for IBM WebSphere Application Server, consider the following scenarios:

A new IBM WebSphere Application Server server machine, hostA.

This scenario starts with adding embedded messaging as the JMS provider, then optionally adding WebSphere MQ as an alternative JMS provider. Each stage summarizes the messaging functions that can be added.

1. Installing embedded messaging as the only JMS provider.

You want to be able to run WebSphere applications that use the WebSphere JMS resources for both point-to-point and publish/subscribe messaging.

- a. Install and customize WebSphere Application Server with the **Embedded Messaging Server** and **Embedded Messaging Client** options. (See the Information Center for details.)
 - b. Use the administrative console to configure WebSphere JMS resources; for example, WebSphere Queue Connection Factories and WebSphere Topic Connection Factories.
 - c. On any client machines that are to use the WebSphere JMS resources, install IBM WebSphere Application Server with the **Embedded Messaging Client** option.
2. Adding WebSphere MQ as an alternative JMS provider for point-to-point messaging.

Besides the point-to-point and publish/subscribe messaging that uses the embedded WebSphere JMS resources (from the preceding step), you want to be use WebSphere MQ Queue resources for point-to-point messaging.

- a. Install WebSphere MQ 5.3 with the required features. For more information about installing WebSphere MQ, see the WebSphere MQ Quick Beginnings book for your machine's platform. When installing WebSphere MQ, perform the following actions:
 - 1) When asked if you want to remove or modify the version of WebSphere MQ currently installed, choose to remove the Server. (The WebSphere MQ install program recognizes the installed WebSphere embedded messaging as an installed version of WebSphere MQ.)
 - 2) Select the option to keep existing queue managers.

- 3) Proceed with the installation of the full WebSphere MQ Server by running the install program again, after the old WebSphere MQ server (the Embedded Messaging Server) has been removed.
 - 4) Do not install the WebSphere MQ Client. Use the WebSphere Application Server **Embedded Messaging Client** that was installed with WebSphere embedded messaging.
 - b. Use the administrative console to configure WebSphere MQ Queue Connection Factories and WebSphere MQ Queue Destinations.
3. Adding WebSphere MQ Event Broker for alternative publish/subscribe messaging.
- For publish/subscribe messaging, you want to be able to run WebSphere applications that use the WebSphere MQ Topic resources or the embedded WebSphere Topic resources (such as those configured in preceding steps).
- a. Install WebSphere MQ Event Broker. For more information about installing WebSphere MQ Event Broker, see the WebSphere MQ Event Broker Installation Guide for your machine's platform.
 - b. Use the administrative console to configure WebSphere MQ Topic Connection Factories and WebSphere MQ Topic Destinations.

An existing WebSphere MQ 5.2 server and broker machine, hostA, where you want to install embedded messaging as the JMS provider.

1. Upgrade to WebSphere MQ 5.3 with the required features. For more information about upgrading to WebSphere MQ 5.3, see "Migrating from an earlier version" in the WebSphere MQ Quick Beginnings book for your machine's platform.
2. To continue using publish/subscribe messaging, upgrade to a supported broker such as WebSphere MQ Event Broker.
3. Install IBM WebSphere Application Server with the **Embedded Messaging Server** and **Embedded Messaging Client** options. (See "Installing the product" in the Information Center.)
4. Use the administrative console to configure WebSphere JMS resources; for example, WebSphere Queue Connection Factories and WebSphere Topic Connection Factories.
5. If you want WebSphere applications to use the WebSphere MQ resources, use the administrative console to configure WebSphere MQ JMS resources; for example, WebSphere MQ Queue Connection Factories and WebSphere MQ Destinations.
6. On any client machines that are to use the WebSphere JMS resources, install IBM WebSphere Application Server with the **Embedded Messaging Client** option. (See "Installing the product" in the Information Center.)

You can run WebSphere applications that use both the WebSphere JMS resources and WebSphere MQ JMS resources for messaging.

An existing WebSphere MQ server machine, hostA, where you want to use WebSphere MQ as the only JMS provider.

1. For point-to-point messaging, ensure that you have installed WebSphere MQ 5.3 with required features. For publish/subscribe messaging, ensure that you have also installed a supported broker such as WebSphere MQ Event Broker. For more information about installing WebSphere MQ 5.3 and Event Broker, see the WebSphere MQ Quick Beginnings book and WebSphere MQ Event Broker Installation Guide for your machine's platform.
2. Install IBM WebSphere Application Server without any of the **Embedded Messaging Server** and **Embedded Messaging Client** options.
3. Use the administrative console to configure WebSphere MQ JMS resources; for example, WebSphere MQ Queue Connection Factories and WebSphere MQ Topic Connection Factories.
4. On any client machines that are to use the WebSphere JMS resources, install IBM WebSphere Application Server with the **Embedded Messaging Client** option.

You can run WebSphere applications that use the WebSphere MQ JMS resources for point-to-point or (with a supported broker installed) publish/subscribe messaging.

Installing WebSphere MQ as the JMS provider:

Use this task to install and configure WebSphere MQ with support for the Java Message Service (JMS) for use with the WebSphere Application Server.

(UNIX platforms only) Before you install WebSphere MQ, create and mount a journalized file system called /var/mqm for your messaging working data. Use a partition strategy with a separate volume for the WebSphere MQ data. This means that other system activity is not affected if a large amount of messaging work builds up in /var/mqm. You can also create separate file systems for your log data (var/mqm/log) and error files (var/mqm/errors). You should store log files on a different physical volume from the messaging queues (var/mqm). This ensures data integrity in the case of a hardware failure. If you are creating separate file systems, allow a minimum of 30 MB of storage for /var/mqm, 20 MB of storage for /var/mqm/log, and 4 MB of storage for /var/mqm/errors.

- The /var file system is used to store all the security logging information for the system, and is used to store the temporary files for email and printing. Therefore, it is critical that you maintain free space in /var for these operations. If you do not create a separate file system for messaging data, and /var fills up, all security logging will be stopped on the system until some free space is available in /var. Also, email and printing will no longer be possible until some free space is available in /var.
- **5.1+** It is not recommended to install WebSphere Studio Application Developer and WebSphere Application Server on the same machine when using WebSphere MQ.
- For more information the security logging issue, see "Security and WebSphere MQ" in the Information Center.
- For more information about creating file systems for WebSphere MQ, and WebSphere MQ space requirements for /var file systems, see the section "Preparing for Installation: Creating WebSphere MQ file systems" in the appropriate WebSphere MQ *Quick Beginnings* book.
- For other installation prerequisites, see the appropriate WebSphere MQ *Quick Beginnings* book, as follows:
 - *WebSphere MQ for Windows, V5.3 Quick Beginnings*, GC34-6073
 - *WebSphere MQ for AIX, V5.3 Quick Beginnings*, GC34-6076
 - *WebSphere MQ for Solaris, V5.3 Quick Beginnings*, GC34-6075
 - *WebSphere MQ for HP-UX, V5.3 Quick Beginnings*, GC34-6077
 - *WebSphere MQ for Linux for Intel and Linux for zSeries, V5.3 Quick Beginnings*, GC34-6078

You can get these books from the WebSphere MQ messaging platform-specific books Web page at <http://www-3.ibm.com/software/ts/mqseries/library/manualsa/manuals/platspecific.html>

To install and configure WebSphere MQ (MQSeries) for use as a JMS provider to IBM WebSphere Application Server, complete the following steps:

1. Install WebSphere MQ 5.3, with the required MQ features, as described in the installation instructions provided with WebSphere MQ.

If you are installing WebSphere MQ on top of WebSphere Application Server embedded messaging, perform the following actions when installing WebSphere MQ:

- a. When asked if you want to remove or modify the version of WebSphere MQ currently installed, choose to remove the Server. (The WebSphere MQ install program recognizes the installed WebSphere embedded messaging as an installed version of WebSphere MQ.)
- b. Select the option to keep existing queue managers.
- c. Proceed with the installation of the full WebSphere MQ Server by running the install program again, after the old WebSphere MQ server (the Embedded Messaging Server) has been removed.
- d. Do not install the WebSphere MQ Client. Use the WebSphere Application Server **Embedded Messaging Client** that was installed with WebSphere embedded messaging.

Note: **5.1+** The WBI Server Foundation package includes copies of the WebSphere MQ 5.3 and Event Broker installation packages, with restricted licensing for use with WBI Server Foundation. (WebSphere MQ Event Broker is not available for Linux.) For more information about the usage restrictions for the WebSphere MQ 5.3 and Event Broker installation packages, see the WBI Server Foundation licensing information.

If you want to use the original WebSphere MQ 5.3 release, ensure that you install the CSD03 update. If you want to use WebSphere MQ 5.3 on the same machine as WebSphere Application Server embedded messaging, ensure that you install the following MQ features:

- For a WebSphere Application Server **Embedded Messaging Server** installation, the required MQ features are “Server” and “Java Messaging”.
- For a WebSphere Application Server **Embedded Messaging Client** installation, the only required MQ feature is “Java Messaging”.

5.1+ You can use the WebSphere MQ 5.3 installation package provided with WBI Server Foundation to install the required MQ features into an existing WebSphere MQ 5.3 or to install a new WebSphere MQ 5.3 with the required MQ features for use with WBI Server Foundation subject to the licensing conditions.

For information about installing WebSphere MQ 5.3, or migrating to WebSphere MQ 5.3 from an earlier release, see the appropriate WebSphere MQ *Quick Beginnings* book, as listed above.

2. If you want to use WebSphere MQ - Publish/Subscribe support, you need to provide a Publish/Subscribe broker.

For example, you can do this by using either WebSphere MQ Event Broker or WebSphere MQ Integrator (formerly MQSeries Integrator). For more information about these products, see the following Web sites:

- WebSphere MQ Event Broker Web site at <http://www-4.ibm.com/software/ts/mqseries/platforms/#eventb>
- WebSphere MQ Integrator Web site at <http://www-4.ibm.com/software/ts/mqseries/platforms/#integrator>

3. Follow the WebSphere MQ 5.3 instructions for verifying your installation setup.
4. For AIX, see the WebSphere MQ 5.3 readme.txt for additional steps.
5. If you want to install IBM WebSphere Application Server on the same host as WebSphere MQ, and have not yet done so, install IBM WebSphere Application Server. If you do not want to use the embedded WebSphere JMS provider, you can install WebSphere MQ then install WebSphere Application Server without the **Embedded Messaging Server** option. You are recommended to install and use the WebSphere Application Server **Embedded Messaging Client**.
6. Set the MQJMS_LIB_ROOT environment variable to the directory where WebSphereMQJava\lib is installed. IBM WebSphere Application Server uses the MQJMS_LIB_ROOT to locate the WebSphere MQ libraries for the WebSphere MQ JMS Provider.

This task has installed WebSphere MQ for use as the JMS provider with WebSphere Application Server.

You can configure JMS resources to be provided by WebSphere MQ, by using the WebSphere administrative console to define WebSphere MQ resources.

(UNIX platforms only) Restrict access to the messaging errors directories and logging files, by using the following commands. This is part of the procedure to secure the directories and log files needed for WebSphere MQ, as described in *Securing messaging directories and log files*.

1. For the /var/mqm/errors directory:

```
chmod 3777 /var/mqm/errors
chown mqm:mqm /var/mqm/errors
```

```
touch /var/mqm/errors/AMQERR01.LOG
chown mqm:mqm /var/mqm/errors/AMQERR01.LOG
chmod 666 /var/mqm/errors/AMQERR01.LOG
```

```
touch /var/mqm/errors/AMQERR02.LOG
chown mqm:mqm /var/mqm/errors/AMQERR02.LOG
chmod 666 /var/mqm/errors/AMQERR02.LOG
```



```
touch /var/mqm/errors/AMQERR03.LOG
chown mqm:mqm /var/mqm/errors/AMQERR03.LOG
chmod 666 /var/mqm/errors/AMQERR03.LOG
```

2. For the `/var/mqm/qmgrs/@SYSTEM/errors` directory:

```
chmod 3777 /var/mqm/qmgrs/@SYSTEM/errors
chown mqm:mqm /var/mqm/qmgrs/@SYSTEM/errors
```

```
touch /var/mqm/qmgrs/@SYSTEM/errors/AMQERR01.LOG
chown mqm:mqm /var/mqm/qmgrs/@SYSTEM/errors/AMQERR01.LOG
chmod 666 /var/mqm/qmgrs/@SYSTEM/errors/AMQERR01.LOG
```

```
touch /var/mqm/qmgrs/@SYSTEM/errors/AMQERR02.LOG
chown mqm:mqm /var/mqm/qmgrs/@SYSTEM/errors/AMQERR02.LOG
chmod 666 /var/mqm/qmgrs/@SYSTEM/errors/AMQERR02.LOG
```

```
touch /var/mqm/qmgrs/@SYSTEM/errors/AMQERR03.LOG
chown mqm:mqm /var/mqm/qmgrs/@SYSTEM/errors/AMQERR03.LOG
chmod 666 /var/mqm/qmgrs/@SYSTEM/errors/AMQERR03.LOG
```

Defining a generic JMS provider:

Use this task to define a new JMS provider to WebSphere Application Server, for use instead of the embedded WebSphere JMS provider or a WebSphere MQ JMS provider.

Before starting this task, you should have installed and configured the JMS provider and its resources by using the tools and information provided with the JMS provider.

To define a new generic JMS provider to WebSphere Application Server, use the administrative console to complete the following steps:

1. In the navigation pane, expand **Resources-> Generic JMS Providers** This displays the existing generic JMS providers in the content pane.
2. To define a new generic JMS provider, click **New** in the content pane. Otherwise, to change the definition of an existing JMS provider, click the JMS provider. This displays the properties used to define the JMS provider in the content pane.
3. Specify appropriate properties for the JMS provider.
4. Click **OK**.
5. To save your configuration, click **Save** on the task bar of the Administrative console window.
6. To have the changed configuration take effect, stop then restart the application server.

Installing WebSphere embedded messaging as the JMS provider:

Use this task to install the Embedded Messaging options of IBM WebSphere Application Server for use as the JMS provider.

Before you install the embedded messaging options of IBM WebSphere Application Server, you must complete the following steps:

1. If you want to install embedded messaging on a machine where you already have WebSphere MQ installed, you must ensure that you have upgraded to WebSphere MQ 5.3 with the required MQ features:
 - a. If you have the original WebSphere MQ 5.3 release installed, ensure that you have applied the CSD04 update.

To determine if your WebSphere MQ 5.3 installation is at the required level, run the **mqver** utility provided by WebSphere MQ. The required level as indicated by mqver is shown below:

```
Name:      WebSphere MQ
Version:   530.4 CSD04
...
```

For more information about CSD04 and other WebSphere MQ product support, see the WebSphere MQ support page for your platform.

- b. Ensure that you have installed the following WebSphere MQ features:
 - For a WebSphere Application Server **Embedded Messaging Server** installation, the required MQ features are “Server” and “Java Messaging”.
 - For a WebSphere Application Server **Embedded Messaging Client** installation, the only required MQ feature is “Java Messaging”.

If you have not installed WebSphere MQ 5.3 with the required MQ features, then installation of IBM WebSphere Application Server Embedded Messaging options fails with prerequisite check errors.

5.1+ The WBI Server Foundation package includes copies of the WebSphere MQ 5.3 and Event Broker installation packages, with restricted licensing. (WebSphere MQ Event Broker is not available for Linux.) You can use the provided packages to install the required MQ features or WebSphere MQ 5.3 for use with WBI Server Foundation.

For information about installing WebSphere MQ 5.3, or migrating to WebSphere MQ 5.3 from an earlier release, see the appropriate WebSphere MQ *Quick Beginnings* book, as follows:

- *WebSphere MQ for Windows, V5.3 Quick Beginnings*, GC34-6073
- *WebSphere MQ for AIX, V5.3 Quick Beginnings*, GC34-6076
- *WebSphere MQ for Solaris, V5.3 Quick Beginnings*, GC34-6075
- *WebSphere MQ for HP-UX, V5.3 Quick Beginnings*, GC34-6077
- *WebSphere MQ for Linux for Intel and Linux for zSeries, V5.3 Quick Beginnings*, GC34-6078

You can get these books from the WebSphere MQ messaging platform-specific books Web page at <http://www-3.ibm.com/software/ts/mqseries/library/manualsa/manuals/platspecific.html>

2. Ensure that there is enough space in the file systems where you want to install the Embedded Messaging options and store associated messaging data. On UNIX platforms, you also need to create the required file systems before installing the embedded messaging feature.
 - **(Windows only)** You can specify the file system into which the Embedded Messaging options are installed. The following table lists the default locations for the base messaging functions and the messaging broker functions (for publish/subscribe messaging). The table also provides figures for the file system sizes on which you can base your own calculations.

Table 1. Installation directory and space needed for embedded messaging - Windows platform

	Base messaging	Messaging broker
Installation directory	C:\Program Files\IBM\WebSphere MQ	C:\Program Files\IBM\WebSphere MQ\WEMPS
Typical space needed	70 MB (server) or 15 MB (client)	45 MB (server)

If you are using the Installation wizard to install IBM WebSphere Application Server, you can specify an install location for the Embedded Messaging options during either a Full or Custom install. If you want to use a silent install for IBM WebSphere Application Server, you can specify an install location for the Embedded Messaging options when you tailor the options response file before issuing the command to install silently. For more information about performing a silent installation (including editing the options response file), see “Installing silently” in the Information Center.

- **(UNIX platforms only)** The file system into which the Embedded Messaging options are installed is fixed. The following table lists the default locations for the base messaging functions and the messaging broker functions (for publish/subscribe messaging). The table also provides figures for the file system sizes on which you can base your own calculations.

Table 2. Installation directory and space needed for embedded messaging - UNIX platforms

	Base code	Broker code	Base data	Broker data
	/usr/mqm	/usr/opt/wemps	/var/mqm	/var/wemps
AIX	40MB (server) or 15MB (client)	80MB (server) or 15MB (client)	8MB (server) or 5MB (client)	5MB (server)

Table 2. Installation directory and space needed for embedded messaging - UNIX platforms (continued)

	Base code	Broker code	Base data	Broker data
	/opt/mqm	/opt/wemps	/var/mqm	/var/wemps
Linux/Intel	40MB (server) or 15MB (client)	100MB (server) or 15MB (client)	8MB (server) or 5MB (client)	5MB (server)
HP-UX	40MB (server) or 15MB (client)	105MB (server) or 15MB (client)	8MB (server) or 5MB (client)	5MB (server)
Solaris	40MB (server) or 15MB (client)	70MB (server) or 15MB (client)	20MB (server) or 15MB (client)	5MB (server)

Before you install WebSphere embedded messaging, create and mount a journalized file system called /var/mqm for your messaging working data. Use a partition strategy with a separate volume for the WebSphere MQ data. This means that other system activity is not affected if a large amount of messaging work builds up in /var/mqm. You can also create separate file systems for your log data (var/mqm/log) and error files (var/mqm/errors). Store log files on a different physical volume from the Embedded Messaging queues (var/mqm). This ensures data integrity in the case of a hardware failure. If you are creating separate file systems, allow a minimum of 30 MB of storage for /var/mqm, 20 MB of storage for /var/mqm/log, and 4 MB of storage for /var/mqm/errors.

Note: The /var file system is used to store all the security logging information for the system, and is used to store the temporary files for email and printing. Therefore, it is critical that you maintain free space in /var for these operations. If you do not create a separate file system for messaging data, and /var fills up, all security logging will be stopped on the system until some free space is available in /var. Also, email and printing will no longer be possible until some free space is available in /var.

You have the same options for creating file systems for Embedded Messaging as you do for WebSphere MQ. For example:

- If you cannot install the Embedded Messaging options in the required file system (for example, if it is too small), you can do one of the following *before* installing the Embedded Messaging options:
 - Create and mount a new file system for the installation directory.
 - Create a new directory anywhere on your machine, and create a symbolic link from the required installation directory to the new directory. For example, on AIX:

```
mkdir /bigdisk/mqm
ln -s /bigdisk/mqm /usr/mqm
```

3. Define the operating system groups and users needed for embedded messaging:

- **(UNIX platforms only)**
 - a. If you have not already done so, create the groups **mqm** and **mqbrkrs**.
 - b. Add the users **mqm** and **root** to the **mqm** group.
 - c. Add the user **root** to the **mqbrkrs** group.

Notes

Note:

- 1) You are recommended to run the JMS server process under the root user ID. If you run the JMS server process under another user ID, add that user ID to the **mqm** and **mqbrkrs** groups
For more information about running servers under a non-root user ID, see "Running an Application Server with a non-root user ID and the nodeagent as root" in the Information Center.
- 2) User IDs longer than 12 characters cannot be used for authentication with the embedded WebSphere JMS provider.

- **(Windows only)** Define the process user ID with these authorizations:
 - Assign the user ID to the Administrator group.

- Give the user ID the advanced user right, Act as part of the operating system.
- Give the user ID the advanced user right, Log on as a service.

User IDs longer than 12 characters cannot be used for authentication with the embedded WebSphere JMS provider. For example, the default Windows NT user ID, Administrator, is not valid for use with embedded WebSphere messaging, because it contains 13 characters.

The IBM WebSphere Application Server installation wizard GUI grants your Windows user ID the advanced user rights, if the user ID belongs to the Administrator group. The silent installation does not. If you create a new user ID on a Windows platform to perform the silent installation, you must restart the system to activate the proper authorizations for the user ID, and to perform a successful silent installation.

This user ID, the WebSphere Application Server process user ID, is used to start the JMS server (for general JMS support) and the WebSphere Embedded Broker (for WebSphere topic connections).

Ensure that these operating system security settings are used when you next start IBM WebSphere Application Server; for example, either log off and then on again with the process user ID, or open a new shell in which to start IBM WebSphere Application Server.

You have the same options for creating user IDs and groups for Embedded Messaging as you do for WebSphere MQ. For more information about creating user IDs and groups for WebSphere MQ, see the section “Preparing for Installation: Setting up the user ID and group” in the appropriate WebSphere MQ *Quick Beginnings* book, as listed above.

4. **(Solaris only)** Several Solaris kernel values are typically too small for the Embedded Messaging options. Starting the internal JMS server or client with insufficient kernel resources produces a First Failure Support Technology (FFST) file in the `/var/mqm/errors` directory.

Before installing Embedded Messaging, review the machine’s configuration. To do this type the following command:

```
sysdef -i
```

The kernel values are set in the `/etc/system` file, as shown in the following example.

```
set shmsys:shminfo_shmmax = 4294967295
set shmsys:shminfo_shmseg = 1024
set shmsys:shminfo_shmmni = 1024
set semsys:seminfo_semaem = 16384
set semsys:seminfo_semmni = 1024
set semsys:seminfo_semmap = 1026
set semsys:seminfo_semmns = 16384
set semsys:seminfo_semmsl = 100
set semsys:seminfo_semopm = 100
set semsys:seminfo_semmnu = 2048
set semsys:seminfo_semume = 256
set msgsys:msginfo_msgmap = 1026
set msgsys:msginfo_msgmax = 65535
set rlim_fd_cur=1024
```

You can change kernel values by editing the `/etc/system` file then rebooting the operating system. For more information about setting up the Solaris system, see the Solaris System Administration documentation; for example, the *Solaris Tunable Parameters Reference Manual*.

Note: Queue managers are generally independent of each other. Therefore system kernel parameters, for example `shmmni`, `semmni`, `semmns`, and `semmnu` need to allow for the number of queue managers in the system.

5. **(HP-UX only)** Several HP-UX kernel values are typically too small for the Embedded Messaging options. Starting the internal JMS server or client with insufficient kernel resources produces a First Failure Support Technology (FFST) file in the `/var/mqm/errors` directory.

Before installing Embedded Messaging, review the machine’s configuration and, if needed, set appropriate HP-UX kernel settings.

You should set new values for the following messaging-related kernel parameters (along with the base set of kernel parameters required by IBM WebSphere Application Server):

Parameter	Value
sema	1
semaem	16384
semmns	16384
semvmx	32767
shmem	1

This table lists only those messaging-related kernel parameters that either are extra to the base set or need values that are greater than given in the base set.

For information about how to review and set the base set of kernel parameters for IBM WebSphere Application Server, see the item “Configure HP-UX kernel settings before installing” in “Platform-specific tips for installing and migrating (HP-UX platforms)” in the Information Center.

Note: Queue managers are generally independent of each other. Therefore system kernel parameters, for example shmmni, semmni, semmns, and semmnu need to allow for the number of queue managers in the system.

To install the Embedded Messaging options of WebSphere Application Server for use as the WebSphere JMS provider, complete the following steps:

1. Login with the WebSphere Application Server process user ID (defined as part of the prerequisites)

Note: (UNIX platforms only) Login as root.

2. On a machine where you want to host queues or topics, install IBM WebSphere Application Server with the **Embedded Messaging Server** option. If you also want application servers on the host to run messaging applications, install the **Embedded Messaging Client** option.

Both options are selected by default.

3. On a machine where you want application servers to run messaging applications that use a JMS provider on another host, install IBM WebSphere Application Server with the **Embedded Messaging Client** option.

This task has installed WebSphere Application Server with its embedded messaging as the JMS provider.

You can configure JMS resources to be provided by embedded messaging, by using the WebSphere administrative console to define WebSphere JMS resources.

(UNIX platforms only) Restrict access to the messaging errors directories and logging files; for example, by using the following commands:

1. For the /var/mqm/errors directory:

```
chmod 3777 /var/mqm/errors
chown mqm:mqm /var/mqm/errors
```

```
touch /var/mqm/errors/AMQERR01.LOG
chown mqm:mqm /var/mqm/errors/AMQERR01.LOG
chmod 666 /var/mqm/errors/AMQERR01.LOG
```

```
touch /var/mqm/errors/AMQERR02.LOG
chown mqm:mqm /var/mqm/errors/AMQERR02.LOG
chmod 666 /var/mqm/errors/AMQERR02.LOG
```

```
touch /var/mqm/errors/AMQERR03.LOG
chown mqm:mqm /var/mqm/errors/AMQERR03.LOG
chmod 666 /var/mqm/errors/AMQERR03.LOG
```

2. For the /var/mqm/qmgrs/@SYSTEM/errors directory:

```

chmod 3777 /var/mqm/qmgrs/@SYSTEM/errors
chown mqm:mqm /var/mqm/qmgrs/@SYSTEM/errors

touch /var/mqm/qmgrs/@SYSTEM/errors/AMQERR01.LOG
chown mqm:mqm /var/mqm/qmgrs/@SYSTEM/errors/AMQERR01.LOG
chmod 666 /var/mqm/qmgrs/@SYSTEM/errors/AMQERR01.LOG

touch /var/mqm/qmgrs/@SYSTEM/errors/AMQERR02.LOG
chown mqm:mqm /var/mqm/qmgrs/@SYSTEM/errors/AMQERR02.LOG
chmod 666 /var/mqm/qmgrs/@SYSTEM/errors/AMQERR02.LOG

touch /var/mqm/qmgrs/@SYSTEM/errors/AMQERR03.LOG
chown mqm:mqm /var/mqm/qmgrs/@SYSTEM/errors/AMQERR03.LOG
chmod 666 /var/mqm/qmgrs/@SYSTEM/errors/AMQERR03.LOG

```

This is part of the procedure to secure the directories and log files needed for WebSphere embedded messaging, as described in *Securing messaging directories and log files*.

If you have installed the Embedded Messaging Server option on top of WebSphere MQ, the MQ command **setmqcap** is set to use parameter 0 instead of -1, which results in:

- Issuing a license-unit message to the MQ console window whenever a queue manager starts
- Writing a message to the MQ error log

To prevent this, after you have completed the installation of IBM WebSphere Application Server, issue the `setmqcap -1` command from a command line.

Moving from the embedded WebSphere JMS provider to WebSphere MQ

Use this task to move from the embedded WebSphere JMS provider to WebSphere MQ as the provider of messaging services and resources for WebSphere enterprise applications.

To move from the embedded WebSphere JMS provider to WebSphere MQ as the provider of messaging services and resources for WebSphere enterprise applications, you need to install and configure a supported level of WebSphere MQ with the required MQ features.

Existing JMS resource definitions for the embedded WebSphere JMS provider continue to work with WebSphere MQ as the JMS provider, so you do not need to redefine those JMS resources. However, to take advantage of the extra configuration options for WebSphere MQ, you can use the administrative console to define new JMS resources as WebSphere MQ resources; for example, to define MQ Queue Connection Factories.

1. For WebSphere MQ point-to-point messaging, install the base WebSphere MQ product.

When installing WebSphere MQ, perform the following actions:

- a. When asked if you want to remove or modify the version of WebSphere MQ currently installed, choose to remove the Server. (The WebSphere MQ install program recognizes the installed WebSphere embedded messaging as an installed version of WebSphere MQ.)
- b. Select the option to keep existing queue managers.
- c. Proceed with the installation of the full WebSphere MQ Server by running the install program again, after the old WebSphere MQ server (the Embedded Messaging Server) has been removed.
- d. Do not install the WebSphere MQ Client. Use the WebSphere Application Server **Embedded Messaging Client** that was installed with WebSphere embedded messaging.

For information about installing WebSphere MQ 5.3, or migrating to WebSphere MQ 5.3 from an earlier release, see the appropriate WebSphere MQ *Quick Beginnings* book, as follows:

- *WebSphere MQ for Windows, V5.3 Quick Beginnings*, GC34-6073
- *WebSphere MQ for AIX, V5.3 Quick Beginnings*, GC34-6076
- *WebSphere MQ for Solaris, V5.3 Quick Beginnings*, GC34-6075
- *WebSphere MQ for Linux for Intel and Linux for zSeries, V5.3 Quick Beginnings*, GC34-6078

You can get WebSphere MQ messaging platform-specific books at this Web page:
<http://www-3.ibm.com/software/ts/mqseries/library/manualsa/manuals/platspecific.html>

2. Configure WebSphere MQ queue resources to IBM WebSphere Application Server.
 - For point-to-point messaging WebSphere applications can continue to use WebSphere queue resources (through the embedded messaging JMS provider) or WebSphere MQ queue resources.
 - For publish/subscribe messaging, WebSphere applications can continue to use WebSphere topic resources (through the embedded messaging JMS provider).
3. For WebSphere MQ publish/subscribe messaging, install a Publish/Subscribe broker, such as WebSphere MQ Event Broker.

For more information about installing WebSphere MQ Event Broker or WebSphere MQ Integrator (formerly MQSeries Integrator), see the following Web sites:

 - WebSphere MQ Event Broker Web site at <http://www-4.ibm.com/software/ts/mqseries/platforms/#eventb>
 - WebSphere MQ Integrator Web site at <http://www-4.ibm.com/software/ts/mqseries/platforms/#integrator>
4. Configure WebSphere MQ topic resources to IBM WebSphere Application Server. For publish/subscribe messaging, WebSphere applications can continue to use WebSphere topic resources (through the embedded messaging JMS provider) or WebSphere MQ topics.

Managing JMS servers on an Application Server node

Use this task to manage JMS servers on an Application Server node that is not part of a deployment manager cell.

On an Application Server node that is not part of a deployment manager cell, each application server has an internal JMS server that runs as part of the application server process and is administered as additional properties of the application server. A JMS server enables the application server to access JMS resources.

If you want to start a JMS server that is not part of a deployment manager cell, you start the associated application server. For more information about starting application servers, see "Starting servers" in the Information Center.

You can use the WebSphere administrative console to configure a general set of JMS server properties, which add to the default values of properties configured automatically for the embedded WebSphere JMS provider.

To configure JMS server properties, use the administrative console to complete the following steps:

1. In the navigation pane, select **Servers-> Application Servers** This displays a table of the application servers in the administrative domain.
2. In the content pane, click the name of the application server. This displays the properties of the application server in the content pane.
3. In the content pane, under Additional Properties, select **Server components-> JMS Server** This displays the JMS server properties in the content pane.
4. Specify appropriate properties for the JMS server.

If you want the JMS server to be started automatically when the application server is next started, set the **Initial state** property to started.

If you want to add a new queue to be hosted by the JMS server, add the administrative name of the queue to the Queue Names field. (The name must match the name of a WebSphere Queue administrative object, including the use of upper- and lowercase.) Similarly, if you want to remove a queue from the JMS server, remove its name from that field.

5. Click **OK**.
6. To save your configuration, click **Save** on the task bar of the administrative console window.
7. To have the changed configuration take effect, stop then restart the application server.

Managing JMS servers in a deployment manager cell

Use this task to manage JMS servers on nodes in a WebSphere Application Server deployment manager cell.

In a WebSphere Application Server deployment manager cell, each node can have at most one JMS server, and any application server within the cell can access JMS resources served by any of those JMS servers.

You can use the WebSphere administrative console to display a list of all JMS servers, to show and control their runtime status. You can also configure a general set of JMS server properties, which add to the default values of WebSphere MQ properties configured automatically for the embedded WebSphere JMS provider.

Note: In general, the default values of WebSphere MQ properties are adequate for WebSphere internal JMS servers. However, if you are running high messaging loads, you may need to change some WebSphere MQ properties; for example, WebSphere MQ properties for log file locations, file pages, and buffer pages. For more information about configuring WebSphere MQ properties, see the *WebSphere MQ System Administration* book, SC33-1873, which is available from the IBM Publications Center or from the WebSphere MQ collection kit, SK2T-0730.

To manage a WebSphere internal JMS server, use the administrative console to complete the following steps:

1. In the navigation pane, select **Servers-> JMS Servers**. This displays a table of the JMS servers, showing their runtime status.
2. If you want to change the runtime status of a JMS server, complete the following steps:
 - a. In the table of JMS servers, select the JMS servers that you want to act on.
 - To act on one or more specific JMS servers, select the checkbox next to the JMS server name.
 - To act on all JMS servers, select the checkbox next to the JMS servers title of the table.
 - b. Click one of the actions displayed to change the status of the JMS servers; for example, click **Stop** to stop a JMS server.

The status of the JMS servers that you have acted on is updated to show the result of your actions.

3. If you want to change the properties of a JMS server, click the name of the JMS server. This displays the properties of the JMS server in the content pane.
4. Specify appropriate properties for the JMS server.

If you want the JMS server to be started automatically, set the **Initial state** property to started.

If you want to add a new queue to be hosted by the JMS server, add the administrative name of the queue to the Queue Names field. (The name must match the name of a WebSphere Queue administrative object, including the use of upper- and lowercase.) Similarly, if you want to remove a queue from the JMS server, remove its name from that field.

5. Click **OK**.
6. To save your configuration, click **Save** on the task bar of the Administrative console window.
7. To have the changed configuration take effect, stop then restart the JMS Server.

Configuring JMS provider resources

Use the following tasks to configure JMS provider resources needed to support enterprise beans that exploit JMS services.

- Configuring resources for the embedded WebSphere JMS provider
 - Configuring a queue connection factory
 - Configuring a topic connection factory
 - Configuring a queue destination
 - Configuring a topic destination
- Configuring resources for the WebSphere MQ JMS provider
 - Configuring a queue connection factory

- Configuring a topic connection factory
- Configuring a queue destination
- Configuring a topic destination
- Configuring resources for a generic JMS provider
 - Configuring a JMS connection factory
 - Configuring a JMS destination

Configuring resources for the embedded WebSphere JMS provider:

Use the following tasks to configure the connection factories and destinations for the embedded WebSphere JMS provider.

You only need to complete these tasks if your WebSphere Application Server environment uses the embedded WebSphere JMS provider to support enterprise applications that use JMS.

- Configuring a queue connection factory
- Configuring a topic connection factory
- Configuring a queue destination
- Configuring a topic destination

Configuring a queue connection factory, embedded WebSphere JMS provider:

Use this task to configure the properties of a queue connection factory for use with the embedded WebSphere JMS provider. This task contains an optional step for you to create a new queue connection factory.

To configure the properties of a queue connection factory for use with the embedded WebSphere JMS provider, use the administrative console to complete the following steps:

1. Display the embedded WebSphere JMS provider.
 - a. In the navigation pane, click **Resources-> WebSphere JMS Provider**
 - b. In the content pane, click the name of the WebSphere JMS provider that you want to work with.

This displays in the content pane a table of properties for the WebSphere JMS provider, including links to the types of JMS resources supported by the JMS provider.
2. In the content pane, under Additional Properties, click **WebSphere Queue Connection Factories** This displays any existing queue connection factories for the WebSphere JMS provider in the content pane.
3. To create a new queue connection factory, click **New** in the content pane. Otherwise, to change the properties of an existing queue connection factory, click one of the connection factories displayed. This displays the properties for the queue connection factory in the content pane.
4. Specify appropriate properties for the queue connection factory.
5. Click **OK**.
6. To save your configuration, click **Save** on the task bar of the Administrative console window.
7. To have the changed configuration take effect, stop then restart the application server.

Configuring a topic connection factory, embedded WebSphere JMS provider:

Use this task to configure the properties of a topic connection factory for use with the embedded WebSphere JMS provider. This task contains an optional step for you to create a new topic connection factory.

To configure the properties of a topic connection factory for use with the embedded WebSphere JMS provider, use the administrative console to complete the following steps:

1. Display the embedded WebSphere JMS provider.
 - a. In the navigation pane, click **Resources-> WebSphere JMS Provider**
 - b. In the content pane, click the name of the WebSphere JMS provider that you want to work with.

This displays in the content pane a table of properties for the WebSphere JMS provider, including links to the types of JMS resources supported by the JMS provider.

2. In the content pane, under Additional Properties, click **WebSphere Topic Connection Factories** This displays any existing topic connection factories for the WebSphere JMS provider in the content pane.
3. To create a new topic connection factory, click **New** in the content pane. Otherwise, to change the properties of an existing topic connection factory, click one of the connection factories displayed. This displays the properties for the topic connection factory in the content pane.
4. Specify appropriate properties for the topic connection factory.
5. Click **OK**.
6. To save your configuration, click **Save** on the task bar of the Administrative console window.
7. To have the changed configuration take effect, stop then restart the application server.

Configuring a queue destination, embedded WebSphere JMS provider:

Use this task to configure the properties of a queue destination for use with the embedded WebSphere JMS provider. This task contains an optional step for you to create a new queue destination.

To configure the properties of a queue destination for use with the embedded WebSphere JMS provider, use the administrative console to complete the following steps:

1. In the navigation pane, expand **Resources-> WebSphere JMS Provider** This displays in the content pane a table of properties for the WebSphere JMS provider, including links to the types of JMS resources supported by the JMS provider.
2. In the content pane, under Additional Properties, click **WebSphere Queue Destinations** This displays any existing queue destinations for the WebSphere JMS provider in the content pane.
3. To create a new queue destination, click **New** in the content pane. Otherwise, to change the properties of an existing queue destination, click one of the destinations displayed. This displays the properties for the queue destination in the content pane.
4. Specify appropriate properties for the queue destination.
5. Click **OK**.
6. To save your configuration, click **Save** on the task bar of the Administrative console window.
7. To make a queue destination available to applications, you need to host the queue on a JMS server. To add a new queue to a JMS server or to change an existing queue on a JMS server, you define the administrative name of the queue to the JMS server, as described in or Managing JMS servers in a deployment manager cell.
8. To have the changed configuration take effect, stop then restart the application server.

Configuring a topic destination, embedded WebSphere JMS provider:

Use this task to configure the properties of a topic destination for use with the embedded WebSphere JMS provider. This task contains an optional step for you to create a new topic destination factory.

To configure the properties of a topic destination factory for use with the embedded WebSphere JMS provider, use the administrative console to complete the following steps:

1. In the navigation pane, expand **Resources-> WebSphere JMS Provider** This displays in the content pane a table of properties for the WebSphere JMS provider, including links to the types of JMS resources supported by the JMS provider.
2. In the content pane, under Additional Properties, click **WebSphere Topic Destinations** This displays any existing topic destinations for the WebSphere JMS provider in the content pane.
3. To create a new topic destination, click **New** in the content pane. Otherwise, to change the properties of an existing topic destination, click one of the destinations displayed. This displays the properties for the topic destination in the content pane.
4. Specify appropriate properties for the topic destination.

5. Click **OK**.
6. To save your configuration, click **Save** on the task bar of the Administrative console window.
7. To have the changed configuration take effect, stop then restart the application server.

Configuring resources for the WebSphere MQ JMS provider:

Use the following tasks to configure the connection factories and destinations for the WebSphere MQ JMS provider.

You only need to complete these tasks if your WebSphere Application Server environment uses the WebSphere MQ JMS provider to support enterprise applications that use JMS. To enable use of the WebSphere MQ JMS provider, you must have installed and configured WebSphere MQ JMS support, as described in *Installing and configuring WebSphere MQ as the JMS provider*.

- Configuring a queue connection factory
- Configuring a topic connection factory
- Configuring a queue destination
- Configuring a topic destination
- Enabling WebSphere MQ JMS connection pooling

Configuring a queue connection factory, WebSphere MQ JMS provider:

Use this task to configure the properties of a queue connection factory for use with the WebSphere MQ JMS provider. This task contains an optional step for you to create a new queue connection factory.

To configure the properties of a queue connection factory for use with the WebSphere MQ JMS provider, use the administrative console to complete the following steps:

1. Display the WebSphere MQ JMS provider.
 - a. In the navigation pane, click **Resources-> WebSphere MQ JMS Provider**
 - b. In the content pane, click the name of the WebSphere MQ JMS provider that you want to work with.

This displays in the content pane a table of properties for the WebSphere MQ JMS provider, including links to the types of JMS resources supported by the JMS provider.

2. In the contents pane, under Additional Properties, click **WebSphere MQ Queue Connection Factories**. This displays a table listing any existing queue connection factories, with a summary of their properties.
3. To create a new queue connection factory, click **New** in the content pane. Otherwise, to change the properties of an existing queue connection factory, click one of the connection factories displayed. This displays the properties for the queue connection factory in the content pane.
4. Specify appropriate properties for the queue connection factory.
5. **5.1+** Optional: Specify any of the following WebSphere MQ Secure Sockets Layer (SSL) properties that you need, as **Custom properties** of the connection factory: SSLPEERNAME, SSLCRL, and SSLCIPHERSUITE.

For more information about these custom properties, see *Custom properties*.

6. Click **OK**.
7. To save your configuration, click **Save** on the taskbar of the Administrative console window.
8. To have the changed configuration take effect, stop then restart the application server.

Configuring a topic connection factory, WebSphere MQ JMS provider:

Use this task to configure the properties of a topic connection factory for use with the WebSphere MQ JMS provider. This task contains an optional step for you to create a new topic connection factory.

To configure the properties of a topic connection factory for use with the WebSphere MQ JMS provider, use the administrative console to complete the following steps:

1. Display the WebSphere MQ JMS provider.
 - a. In the navigation pane, click **Resources-> WebSphere MQ JMS Provider**
 - b. In the content pane, click the name of the WebSphere MQ JMS provider that you want to work with.

This displays in the content pane a table of properties for the WebSphere MQ JMS provider, including links to the types of JMS resources supported by the JMS provider.

2. In the contents pane, under Additional Properties, click **WebSphere MQ Topic Connection Factories**. This displays a table listing any existing topic connection factories, with a summary of their properties.
3. To create a new topic connection factory, click **New** in the content pane. Otherwise, to change the properties of an existing topic connection factory, click one of the connection factories displayed. This displays the properties for the topic connection factory in the content pane.
4. Specify appropriate properties for the topic connection factory.
5. **5.1+** Optional: Specify any of the following WebSphere MQ properties that you need, as **Custom properties** of the connection factory: SSLPEERNAME, SSLCRL, SSLCIPHERSUITE, MSGSELECTION, and SUBSTORE.

For more information about these custom properties, see Custom properties.

6. Click **OK**.
7. To save your configuration, click **Save** on the taskbar of the Administrative console window.
8. To have the changed configuration take effect, stop then restart the application server.

Configuring a queue destination, WebSphere MQ JMS provider:

Use this task to configure the properties of a queue destination for use with the WebSphere MQ JMS provider. This task contains an optional step for you to create a new queue destination factory.

To configure the properties of a queue destination factory for use with the WebSphere MQ JMS provider, use the administrative console to complete the following steps:

1. In the navigation pane, click **Resources-> WebSphere MQ JMS Provider**. This displays in the content pane a table of properties for the WebSphere MQ JMS provider, including links to the types of JMS resources supported by the JMS provider.
2. In the contents pane, under Additional Properties, click **WebSphere MQ Queue Destinations**. This displays a table listing any existing queue destinations, with a summary of their properties.
3. To define a new queue destination, click **New** in the content pane. Otherwise, to change the properties of an existing queue destination, click one of the destinations displayed. This displays the properties for the queue destination in the content pane.
4. Configure appropriate properties for the queue destination.
5. Optional: If you want WebSphere Application Server to try to use the WebSphere MQ queue manager's remote administration utilities to create the queue, configure the WebSphere MQ Queue Connection properties.

If you have already created your underlying queue in WebSphere MQ using its administration tools (such as runmqsc or MQ Explorer), you do not need to configure any of the WebSphere MQ Queue Connection properties. You only need to configure these properties if you want WebSphere Application Server to try to use the WebSphere MQ queue manager's remote administration utilities to create the queue.

Note: For any changes to these properties to take effect on the queue manager, the WebSphere MQ Queue Manager on which the queue resides (or will reside) must be configured for remote administration and be running.

For more details about these properties, see WebSphere MQ config properties for the queue destination.

6. Click **Apply**.
7. To save your configuration, click **Save** on the task bar of the Administrative console window.
8. To have the changed configuration take effect, stop then restart the application server.

Configuring a topic destination, WebSphere MQ JMS provider:

Use this task to configure the properties of a topic destination for use with the WebSphere MQ JMS provider. This task contains an optional step for you to create a new topic destination factory.

To configure the properties of a topic destination factory for use with the WebSphere MQ JMS provider, use the administrative console to complete the following steps:

1. In the navigation pane, click **Resources-> WebSphere MQ JMS Provider** This displays in the content pane a table of properties for the WebSphere MQ JMS provider, including links to the types of JMS resources supported by the JMS provider.
2. In the content pane, under Additional Properties, click **WebSphere MQ Topic Destinations** This displays a table listing any existing topic destinations, with a summary of their properties.
3. To create a new topic destination, click **New** in the content pane. Otherwise, to change the properties of an existing topic destination, click one of the destinations displayed. This displays the properties for the topic destination in the content pane.
4. Specify appropriate properties for the topic destination.
5. Click **OK**.
6. To save your configuration, click **Save** on the task bar of the Administrative console window.
7. To have the changed configuration take effect, stop then restart the application server.

Configuring WebSphere MQ JMS connection pooling:

Use this task to configure properties of WebSphere MQ JMS connection pooling.

To enable WebSphere MQ JMS connection pooling, complete the following steps:

1. Start the WebSphere Administrative console.
2. In the navigation pane, select **Servers-> Application Servers-> your_app_server** This displays the properties of the application server, *your_app_server*, in the content pane.
3. In the Additional Properties table, select **Message Listener Service properties** This displays the Message Listener Service properties in the content pane.
4. Select Custom Properties, then add the following properties:
mjms.pooling.threshold
The maximum number of unused connections in the pool.
mjms.pooling.timeout
The timeout in milliseconds for unused connections in the pool.
5. Click **OK**.
6. To save your configuration, click **Save** on the task bar of the administrative console window.
7. To have the changed configuration take effect, stop then restart the application server.

Configuring resources for a generic JMS provider:

Use the following tasks to configure the connection factories and destinations for a generic JMS provider (not the embedded WebSphere JMS provider or the WebSphere MQ JMS provider).

You only need to complete these tasks if your WebSphere Application Server environment uses another JMS provider to support enterprise applications that use JMS. To enable use of another JMS provider, you must have installed and configured the JMS provider, as described in *Defining a new JMS provider to WebSphere Application Server*.

- Configuring a JMS connection factory
- Configuring a JMS destination

Configuring a JMS connection factory, generic JMS provider:

Use this task to configure the properties of a JMS connection factory for use with a generic JMS provider other than the embedded WebSphere JMS provider or WebSphere MQ.

To configure the properties of a JMS connection factory for use with a generic JMS provider, use the administrative console to complete the following steps:

1. Display the JMS provider.
 - a. In the navigation pane, click **Resources-> Generic JMS Providers**
 - b. In the content pane, click the name of the JMS provider that you want to work with.

This displays in the content pane a table of properties for the JMS provider, including links to the types of JMS resources supported by the JMS provider.
2. In the Additional Properties list in the contents pane, select **JMS Connection Factories** This displays a table listing any existing JMS connection factories, with a summary of their properties.
3. To create a new JMS connection factory, click **New** in the content pane. Otherwise, to change the properties of an existing JMS connection factory, click one of the connection factories displayed. This displays the properties for the JMS connection factory in the content pane.
4. Specify appropriate properties for the JMS connection factory.
5. Click **OK**.
6. To save your configuration, click **Save** on the task bar of the Administrative console window.
7. To have the changed configuration take effect, stop then restart the application server.

Configuring a JMS destination, a generic JMS provider:

Use this task to configure the properties of a JMS destination for use with a generic JMS provider (other than the embedded WebSphere JMS provider or the WebSphere MQ JMS provider).

To configure the properties of a JMS destination for use with a generic JMS provider, use the administrative console to complete the following steps:

1. In the navigation pane, click **Resources-> Generic JMS Providers** This displays in the content pane a list of any existing generic JMS providers.
2. In the content pane, click the JMS provider that you want to support the JMS destination. This displays in the content pane a table of properties for the JMS provider, including links to the types of JMS resources supported by the JMS provider.
3. In the Additional Properties list in the contents pane, select **JMS Destinations** This displays a table listing any existing JMS destinations, with a summary of their properties.
4. To create a new JMS destination, click **New** in the content pane. Otherwise, to change the properties of an existing queue destination, click one of the destinations displayed. This displays the properties for the JMS destination in the content pane.
5. Specify appropriate properties for the JMS destination.
6. Click **OK**.
7. To save your configuration, click **Save** on the task bar of the Administrative console window.
8. To have the changed configuration take effect, stop then restart the application server.

Securing messaging directories and log files

Use this task to restrict access to the /var/mqm directories and log files needed for WebSphere embedded messaging or WebSphere MQ as the JMS provider.

You need to set the permissions described in this topic, to reduce the risk of severe security exposures.

Note: The /var file system is used to store all the security logging information for the system, and is used to store the temporary files for email and printing. Therefore, it is critical that you maintain free space in /var for these operations and prevent unauthorized access to the file system. If you do not create a separate file system for messaging data, and /var fills up, all security logging will be stopped on the system until some free space is available in /var. Also, email and printing will no longer be possible until some free space is available in /var.

This procedure involves steps that you complete at different stages of installing and using IBM WebSphere Application Server, as described below. The steps are also described at appropriate points in other tasks, but are collected here for completeness.

This procedure applies only to the ordinary UNIX file system. If your site uses access-control lists, secure the files by using that mechanism. Any site-specific requirements can affect the desired owner, group and corresponding privileges. For example, on AIX, complete the following steps:

1. Before installing WebSphere embedded messaging or WebSphere MQ, create and mount a journaled file system called /var/mqm. Use a partition strategy with a separate volume for the messaging data. This means that other system activity is not affected if a large amount of messaging work builds up in /var/mqm.
2. Install WebSphere embedded messaging or WebSphere MQ as the JMS provider.

This stage is described in the following topics:

- Installing WebSphere embedded messaging as the JMS provider
- Installing WebSphere MQ as the JMS provider

As part of this stage, the installation program creates the /var/mqm/errors and /var/mqm/qmgrs/@SYSTEM/errors directories used to hold messaging logging files.

3. Restrict access to the /var/mqm/errors directory and the logging files, by using the following commands:

```
chmod 3777 /var/mqm/errors
chown mqm:mqm /var/mqm/errors
```

```
touch /var/mqm/errors/AMQERR01.LOG
chown mqm:mqm /var/mqm/errors/AMQERR01.LOG
chmod 666 /var/mqm/errors/AMQERR01.LOG
```

```
touch /var/mqm/errors/AMQERR02.LOG
chown mqm:mqm /var/mqm/errors/AMQERR02.LOG
chmod 666 /var/mqm/errors/AMQERR02.LOG
```

```
touch /var/mqm/errors/AMQERR03.LOG
chown mqm:mqm /var/mqm/errors/AMQERR03.LOG
chmod 666 /var/mqm/errors/AMQERR03.LOG
```

4. Restrict access to the /var/mqm/qmgrs/@SYSTEM/errors directory and the logging files, by using the following commands:

```
chmod 3777 /var/mqm/qmgrs/@SYSTEM/errors
chown mqm:mqm /var/mqm/qmgrs/@SYSTEM/errors
```

```
touch /var/mqm/qmgrs/@SYSTEM/errors/AMQERR01.LOG
chown mqm:mqm /var/mqm/qmgrs/@SYSTEM/errors/AMQERR01.LOG
chmod 666 /var/mqm/qmgrs/@SYSTEM/errors/AMQERR01.LOG
```

```
touch /var/mqm/qmgrs/@SYSTEM/errors/AMQERR02.LOG
chown mqm:mqm /var/mqm/qmgrs/@SYSTEM/errors/AMQERR02.LOG
chmod 666 /var/mqm/qmgrs/@SYSTEM/errors/AMQERR02.LOG
```



```
touch /var/mqm/qmgrs/@SYSTEM/errors/AMQERR03.LOG
chown mqm:mqm /var/mqm/qmgrs/@SYSTEM/errors/AMQERR03.LOG
chmod 666 /var/mqm/qmgrs/@SYSTEM/errors/AMQERR03.LOG
```

5. For each application server that uses WebSphere embedded messaging or WebSphere MQ, restrict access to the server's `/var/mqm/qmgrs/long_server_name/errors` directory and its messaging logging files. You should restrict access to the server's directory and logging files, as soon after creating the application server as possible.

To restrict access to the server's directory and logging files, use the following commands:

```
chmod 3775 /var/mqm/qmgrs/long_server_name/errors
chown mqm:mqm /var/mqm/qmgrs/long_server_name/errors
```

```
touch /var/mqm/qmgrs/long_server_name/errors/AMQERR01.LOG
chown mqm:mqm /var/mqm/qmgrs/long_server_name/errors/AMQERR01.LOG
chmod 666 /var/mqm/qmgrs/long_server_name/errors/AMQERR01.LOG
```

```
touch /var/mqm/qmgrs/long_server_name/errors/AMQERR02.LOG
chown mqm:mqm /var/mqm/qmgrs/long_server_name/errors/AMQERR02.LOG
chmod 666 /var/mqm/qmgrs/long_server_name/errors/AMQERR02.LOG
```

```
touch /var/mqm/qmgrs/long_server_name/errors/AMQERR03.LOG
chown mqm:mqm /var/mqm/qmgrs/long_server_name/errors/AMQERR03.LOG
chmod 666 /var/mqm/qmgrs/long_server_name/errors/AMQERR03.LOG
```

Where `long_server_name` is the long name assigned to the server, in the following form:

`WAS_nodename_server_name`. For example, if you created an application server called `server1` to run on the node called `apnode1`, the long server name would be: `WAS_apnode1_server1`.

This task has restricted access to the `/var/mqm` directories and log files needed for WebSphere embedded messaging or WebSphere MQ as the JMS provider, such that only the user ID `mqm` or members of the `mqm` user group have write access.

Configuring authorization security for the embedded WebSphere JMS provider

Use this task to configure authorization security for the embedded WebSphere JMS provider.

To configure authorization security for the embedded WebSphere JMS provider complete the following steps.

Security for the embedded WebSphere JMS provider is enabled when you enable global security for WebSphere Application Server. For more information about enabling global security, see "Managing secured applications" in the Information Center.

1. Configure authorization settings to access JMS resources owned by the embedded WebSphere JMS provider. Authorization to access JMS resources owned by the embedded WebSphere JMS provider is controlled by settings in the `WAS_install_root\config\cells\your_cell_name\integral-jms-authorizations.xml` file.

The settings grant or deny authenticated userids access to internal JMS provider resources (queues or topics). As supplied, the `integral-jms-authorisations.xml` file grants the following permissions:

- Read and write permissions to all queues.
- Pub, sub, and persist to all topics.

To configure authorization settings, edit the `integral-jms-authorisations.xml` file according to the information in this topic and in that file. Please note the file is in Unicode, which requires a binary FTP to the host from a workstation.

2. Edit the `queue-admin-userids` element to create a list of userids with administrative access to all queues. Administrative access is needed to create queues and perform other administrative activities on queues. For example, consider the following `queue-admin-userids` section:


```
<queue-admin-userids>
  <userid>adminid1</userid>
  <userid>adminid2</userid>
</queue-admin-userids>
```

In this example the userids adminid1 and adminid2 are defined to have administrative access to all queues.

3. Edit the queue-default-permissions element to define the default queue access permissions. These permissions are used for queues for which you do not define specific permissions (in queue sections). If this section is not specified, then access permissions exist only for those queues for which you have specifically created queue elements.

For example, consider the following queue-default-permissions element:

```
<queue-default-permissions>
  <permission>write</permission>
</queue-default-permissions>
```

In this example the default access permission for all queues is **write**. This can be overridden for a specific queue by creating a queue element that sets its access permission to **read**.

4. If you want to define specific access permissions for a queue, create a queue element, then define the following elements:

For example, consider the following queue element:

```
<queue>
  <name>q1</name>
  <public>
  </public>
  <authorize>
    <userid>useridr</userid>
    <permission>read</permission>
  </authorize>
  <authorize>
    <userid>useridw</userid>
    <permission>write</permission>
  </authorize>
  <authorize>
    <userid>useridrw</userid>
    <permission>read</permission>
    <permission>write</permission>
  </authorize>
</queue>
```

In this example for the queue q1, the userid useridr has read permission, the userid useridw has write permission, the userid useridrw has both read and write permissions, and all other userids have no access permissions (<public></public>).

5. Edit topic elements to define the access permissions for publish/subscribe topic destinations. For topics, you can grant and deny access permissions. Full permission inheritance is supported on topics. If you do not define specific access permissions for a userid on a specific topic then permissions are inherited first from the public permissions on that topic then from the parent topic. The inheritance of access permissions continues until the root topic from which the root permissions are assumed.
 - a. If you want to define default access permissions for the root topic, edit a topic element with an empty name element. If you omit such a topic section, topics have no default topic permissions other than those defined by specific topic elements. For example, consider the following topic element for the root topic:

```
<topic>
  <name></name>
  <public>
    <permission>+pub</permission>
  </public>
</topic>
```

In this example, the default access permission for all topics is set to publish. This can be overridden by other topic elements for specific topic names.

- b. If you want to define access permissions for a specific topic, create a topic element with the name for the topic then define the access permissions in the public and authorize elements of the topic element. For example, consider the following topic section:

```
<topic>
  <name>a/b/c</name>
  <public>
    <permission>+sub</permission>
  </public>
  <authorize>
    <userid>useridpub</userid>
    <permission>+pub</permission>
  </authorize>
</topic>
```

In this example, the subscribe permission is granted to anyone accessing any topic whose name starts with a/b/c. Also, the userid `useridpub` is granted publish permission for any topic whose name starts with a/b/c.

6. Save the `integral-jms-authorizations.xml` file.

If the dynamic update setting is selected, changes to the `integral-jms-authorizations.xml` file become active when the changed file is saved, so there is no need to stop and restart the JMS server. If the dynamic update setting is not selected, you need to stop and restart the JMS server to make changes active.

Authorization settings for embedded WebSphere JMS provider resources:

Use the `integral-jms-authorisations.xml` file to view or change the authorization settings for JMS resources owned by the embedded WebSphere JMS provider.

Authorization to access JMS resources owned by the embedded WebSphere JMS provider is controlled by the following settings in the `was_install\config\cells\your_cell_name\integral-jms-authorisations.xml` file.

This structure of the settings in `integral-jms-authorisations.xml` is shown in the following example. Descriptions of these settings are provided after the example. To configure authorization settings, follow the instructions provided in *Configuring authorization security for the embedded WebSphere JMS provider*

```
<integral-jms-authorizations>
```

```
  <dynamic-update>true</dynamic-update>
```

```
  <queue-admin-userids>
    <userid>adminid1</userid>
    <userid>adminid2</userid>
  </queue-admin-userids>
```

```
  <queue-default-permissions>
    <permission>write</permission>
  </queue-default-permissions>
```

```
  <queue>
    <name>q1</name>
    <public>
    </public>
    <authorize>
      <userid>useridr</userid>
      <permission>read</permission>
    </authorize>
    <authorize>
      <userid>useridw</userid>
      <permission>write</permission>
```

```

    </authorize>
</queue>

<queue>
  <name>q2</name>
  <public>
    <permission>write</permission>
  </public>
  <authorize>
    <userid>useridr</userid>
    <permission>read</permission>
  </authorize>
</queue>

<topic>
  <name></name>
  <public>
    <permission>+pub</permission>
  </public>
</topic>

<topic>
  <name>a/b/c</name>
  <public>
    <permission>+sub</permission>
  </public>
  <authorize>
    <userid>useridpub</userid>
    <permission>+pub</permission>
  </authorize>
</topic>

</integral-jms-authorizations>

```

dynamic-update: Controls whether or not the JMS Server checks dynamically for updates to this file.

true (Default) Enables dynamic update support.

false Disables dynamic update checking and improves authorization performance.

queue-admin-userids: This element lists those userids with administrative access to all WebSphere queue destinations. Administrative access is needed to create queues and perform other administrative activities on queues. You define each userid within a separate userid sub element:

<userid>adminid</userid>

Where *adminid* is a user ID that can be authenticated by IBM WebSphere Application Server.

queue-default-permissions: This element defines the default queue access permissions that are assumed if no permissions are specified for a specific queue name. These permissions are used for queues for which you do not define specific permissions (in queue elements). If this element is not specified, then no access permissions exist unless explicitly authorized for individual queues.

You define the default permission within a separate permission sub element:

<permission>read-write</permission>

Where *read-write* is one of the following keywords:

read By default, userids have read access to WebSphere queue destinations.

write By default, userids have write access to WebSphere queue destinations.

queue: This element contains the following authorization settings for a single queue destination:

name The name of the queue.

public The default public access permissions for the queue. This is used only for those userids that have no specific authorize element. If you leave this element empty, or do not define it at all, only those userids with authorize elements can access the queue.

You define each default permission within a separate permission element.

authorize

The access permissions for a specific userid. Within each authorize element, you define the following elements:

userid The userid that you want to assign a specific access permission.

permission

An access permission for the associated userid.

You define each permission within a separate permission element. Each permission element can contain the keyword read or write to define the access permission.

For example, consider the following queue element:

```
<queue>
  <name>q1</name>
  <public>
</public>
  <authorize>
    <userid>useridr</userid>
    <permission>read</permission>
  </authorize>
  <authorize>
    <userid>useridw</userid>
    <permission>write</permission>
  </authorize>
  <authorize>
    <userid>useridrw</userid>
    <permission>read</permission>
    <permission>write</permission>
  </authorize>
</queue>
```

topic: This element contains the following authorization settings for a single topic destination:

Each topic element has the following sub elements:

name The name of the topic, without wildcards or other substitution characters.

public The default public access permissions for the topic. This is used only for those userids that have no specific authorize element. If you leave this element empty, or do not define it at all, only those userids with authorize elements can access the topic.

You define each default permission within a separate permission element.

authorize

The access permissions for a specific userid. Within each authorize element, you define the following elements:

userid The userid that you want to assign a specific access permission.

permission

An access permission for the associated userid.

You define each permission within a separate permission element. Each permission element can contain one of the following keywords to define the access permission:

+pub Grant publish permission

+sub Grant subscribe permission

+persist

Grant persist permission

-pub Deny publish permission

-sub Deny subscribe permission

-persist

Deny persist permission

Displaying administrative lists of JMS resources

Use this task with the WebSphere administrative console to display administrative lists of JMS resources.

You can use the WebSphere administrative console to display lists of the following types of JMS resources. You can use the panels displayed to select JMS resources to administer, or to create or delete JMS resources (where appropriate).

To display administrative lists of JMS resources, complete the following general steps:

1. Start the WebSphere administrative console.
2. In the navigation pane, expand the appropriate path to select the type of JMS provider (as shown in the following table).
3. If appropriate, in the content pane, select a specific JMS provider. This displays the properties for the JMS provider, and an Additional Properties list of links to the types of JMS resources provided.
4. In the content pane, under Additional Resources, select the link for the type of JMS resource. This displays a list of the selected JMS resource type in the content pane.

Table 3. Network Deployment and Enterprise - Administrative panels for JMS resources

Path	Panel	Description
Servers-> JMS servers	JMS servers	List all JMS servers within the administration domain
Servers-> JMS servers-> <i>server_name</i>	JMS servers	List properties of the selected JMS server <i>server_name</i>
Embedded WebSphere JMS providers		
Resources-> WebSphere JMS Provider	WebSphere JMS providers	List all WebSphere JMS providers within the administration domain
Resources-> WebSphere JMS Providers-> (In content pane) <i>WebSphereJMSProvider</i>	WebSphere JMS provider	List properties and resources of the selected WebSphere JMS provider
Resources-> WebSphere JMS Providers-> (In content pane) <i>WebSphereJMSProvider</i> -> (Under Additional Properties) WebSphere JMS Queue Connection Factories	WebSphere queue connection factories	List all queue connection factories of the selected WebSphere JMS provider
Resources-> WebSphere JMS Providers-> (In content pane) <i>WebSphereJMSProvider</i> -> (Under Additional Properties) WebSphere Topic Connection Factories	WebSphere topic connection factories	List all topic connection factories of the selected WebSphere JMS provider
Resources-> WebSphere JMS Providers-> (In content pane) <i>WebSphereJMSProvider</i> -> (Under Additional Properties) WebSphere Queue Destinations	WebSphere queue destinations	List all queue destinations of the selected WebSphere JMS provider
Resources-> WebSphere JMS Providers-> (In content pane) <i>WebSphereJMSProvider</i> -> (Under Additional Properties) WebSphere Topic Destinations	WebSphere topic destinations	List all topic destinations of the selected WebSphere JMS provider
WebSphere MQ JMS provider		
Resources-> WebSphere MQ JMS Providers	WebSphere MQ JMS providers	List properties and resource types for the WebSphere MQ JMS provider
Resources-> WebSphere MQ JMS Providers-> (In content pane) <i>MQJMSProvider</i>	WebSphere MQ JMS provider	List properties and resources of the selected WebSphere MQ JMS provider

Table 3. Network Deployment and Enterprise - Administrative panels for JMS resources (continued)

Path	Panel	Description
Resources-> WebSphere MQ JMS Providers-> (In content pane) <i>MQJMSProvider</i> -> (Under Additional Properties) WebSphere MQ Queue connection factories	WebSphere MQ queue connection factories	List all queue connection factories of the selected WebSphere MQ JMS provider
Resources-> WebSphere MQ JMS providers-> (In content pane) <i>MQJMSProvider</i> -> (Under Additional Properties) WebSphere MQ Topic connection factories	WebSphere MQ topic connection factories	List all topic connection factories of the selected WebSphere MQ JMS provider
Resources-> WebSphere MQ JMS Providers-> (In content pane) <i>MQJMSProvider</i> -> (Under Additional Properties) WebSphere MQ Queue destinations	WebSphere MQ queue destination	List all queue destinations of the selected WebSphere MQ JMS provider
Resources-> WebSphere MQ JMS providers-> (In content pane) <i>MQJMSProvider</i> -> WebSphere MQ Topic destinations	WebSphere MQ topic destination	List all topic destinations of the selected WebSphere MQ JMS provider
Generic JMS providers		
A JMS provider other than the embedded WebSphere JMS provider or the WebSphere MQ JMS provider		
Resources-> Generic JMS Providers	Generic JMS providers	List all generic JMS providers within the administration domain
Resources-> Generic JMS Providers-> (In content pane) <i>provider_name</i>	Generic JMS provider <i>provider_name</i>	List properties and resources of the selected generic JMS provider <i>provider_name</i>
Resources-> Generic JMS Providers-> (In content pane) <i>provider_name</i> ->(Under Additional Properties) JMS connection factories	Generic JMS connection factories	List all JMS connection factories of the selected generic JMS provider <i>provider_name</i>
Resources-> Generic JMS Providers-> (In content pane) <i>provider_name</i> -> (Under Additional Properties) JMS destinations	Generic JMS destinations	List all JMS destinations (queues and topics) of the selected generic JMS provider <i>provider_name</i>

JMS server collection:

Each JMS server provides the functions of the JMS provider for a node in your administrative domain. Use this panel to list the JMS servers within the administration domain, or to select a JMS server to view or change its configuration properties.

There can be at most one JMS server on each node in the administration domain, and any application server within the domain can access JMS resources served by any JMS server on any node in the domain.

To view this administrative console page, use the administrative console to complete the following steps:

1. In the navigation pane, select **Servers-> JMS Servers**.

To view or change the properties of a JMS server, select its name in the list displayed.

To act on one or more of the JMS servers listed, click the check boxes next to the names of the objects that you want to act on, then use the buttons provided.

JMS provider collection:

Use this panel to list JMS providers, or to select a JMS provider to view or change its configuration properties.

To view this administrative console page, use the administrative console to complete the following steps:

1. In the navigation pane, expand one of the following paths:

- **Resources-> WebSphere JMS Provider**
- **Resources-> WebSphere MQ JMS Provider**
- **Resources-> Generic JMS Providers**

To view or change the properties of a JMS provider or its resources, select its name in the list displayed.

To define a new generic JMS provider, on the **Resources-> Generic JMS Providers** page click **New**.

To act on one or more of the JMS providers listed, click the check boxes next to the names of the objects that you want to act on, then use the buttons provided.

Name The name by which this JMS provider is known for administrative purposes.

Description

A description of this JMS provider for administrative purposes.

WebSphere JMS provider settings:

Use this panel to view the configuration properties of the embedded WebSphere JMS provider that is installed with WebSphere Application Server. *You cannot change these properties.*

To view this administrative console page, click **Resources-> WebSphere JMS Provider**

Name:

The name by which the JMS provider is known for administrative purposes.

Data type	String
Default	WebSphereJMSPProvider

Description:

A description of the JMS provider, for administrative purposes within IBM WebSphere Application Server.

Data type	String
Default	Built-in WebSphere JMS Provider

WebSphere MQ JMS provider settings:

Use this panel to view the configuration properties of the WebSphere MQ JMS provider. These properties apply only if you have installed WebSphere MQ as the JMS provider over the internal JMS provider installed with WebSphere Application Server. *You cannot change these properties.*

To view this administrative console page, click **Resources-> WebSphere MQ JMS Provider**

Name:

The name by which the WebSphere MQ JMS provider is known for administrative purposes.

Data type	String
Default	WebSphereMQJMSProvider

Description:

A description of the JMS provider, for administrative purposes within IBM WebSphere Application Server.

Data type	String
Default	WebSphere MQ JMS provider

Classpath:

The Java classpath for the JMS provider. The list of paths or JAR file names that together form the location for the JMS provider classes.

Data type	String
Default	\$MQJMS_LIB_ROOT

Native Library Path:

The native library path for the JMS provider. An optional path to any native libraries needed by the JMS provider.

Data type	String
Default	\$MQJMS_LIB_ROOT

The Native Library Path property is set to the directory where the WebSphere MQ Java feature is installed.

JMS provider settings:

If you want to use a JMS provider other than the embedded WebSphere JMS provider or the WebSphere MQ JMS provider, use this panel to configure properties of the JMS provider.

To view this administrative console page, click **Resources-> Generic JMS Providers-> provider_name**

Name:

The name by which the JMS provider is known for administrative purposes.

Data type	String
Default	Null

Description:

A description of the JMS provider, for administrative purposes

Data type	String
Default	Null

Classpath:

The Java classpath for the JMS provider. The list of paths or JAR file names that together form the location for the JMS provider classes.

Data type	String
Default	Null

Native Library Path:

The native library path for the JMS provider. An optional path to any native libraries needed by the JMS provider.

Data type	String
Default	Null

External initial context factory:

The Java classname of the initial context factory for the JMS provider.

For example, for an LDAP service provider the value has the form: com.sun.jndi.ldap.LdapCtxFactory.

Data type	String
Default	Null

External provider URL:

The JMS provider URL for external JNDI lookups.

For example, an LDAP URL for a JMS provider has the form: ldap://hostname.company.com/contextName.

Data type	String
Default	Null

WebSphere Queue connection factory collection:

The queue connection factories configured in the embedded WebSphere JMS provider for point-to-point messaging with JMS queues.

This panel shows a list of the WebSphere queue connection factories with a summary of their configuration properties.

To view this administrative console page, use the administrative console to complete the following steps:

1. In the navigation pane, expand **Resources-> WebSphere JMS Provider**.
2. In the content pane, select the JMS provider that you want to support the connection factory.
3. In the Additional Properties list in the contents pane, select **WebSphere Queue Connection Factory**.

To view or change the properties of a connection factory, select its name in the list displayed.

To act on one or more of the connection factories listed, click the check boxes next to the names of the objects that you want to act on, then use the buttons provided.

WebSphere queue connection factory settings:

Use this panel to view or change the configuration properties of the selected queue connection factory for use with the embedded WebSphere JMS provider that is installed with WebSphere Application Server. These configuration properties control how connections are created to the associated JMS queue destination.

A queue connection factory is used to create JMS connections to queue destinations. The queue connection factory is created by the embedded WebSphere JMS provider. A queue connection factory for the embedded WebSphere JMS provider has the following properties:

To view this administrative console page, click **Resources-> WebSphere JMS Provider-> (In content pane, under Additional Properties) WebSphere Queue Connection Factories-> *connection_factory***

Scope:

Specifies the level to which this resource definition is visible -- the cell, node, or server level.

Resources such as JMS Providers, Namespace bindings, or shared libraries can be defined at multiple scopes, with resources defined at more specific scopes overriding duplicates which are defined at more general scopes.

Note that no matter what the scope of a defined resource, the resource's properties only apply at an individual server level. For example, if you define the scope of a data source at the Cell level, all users in that Cell can look up and use that data source, which is unique within that Cell. However, resource property settings are local to each server in the Cell. For example, if you set **Max Connections** to 10, then each server in that Cell can have 10 connections.

Cell The most general scope. Resources defined at the Cell scope are visible from all Nodes and servers, unless they are overridden. To view resources defined in the cell scope, do not specify a server or a node name in the scope selection form.

Node The default scope for most resource types. Resources defined at the Node scope override any duplicates defined at the Cell scope and are visible to all servers on the same node, unless they are overridden at a server scope on that node. To view resources defined in a node scope, do not specify a server, but select a node name in the scope selection form.

Server The most specific scope for defining resources. Resources defined at the Server scope override any duplicate resource definitions defined at the Cell scope or parent Node scope and are visible only to a specific server. To view resources defined in a server scope, specify a server name as well as a node name in the scope selection form.

When resources are created, they are always created into the current scope selected in the panel. To view resources in other scopes, specify a different node or server in the scope selection form.

Data type String

Name:

The name by which this queue connection factory is known for administrative purposes within IBM WebSphere Application Server. The name must be unique within the JMS connection factories across the WebSphere administrative domain.

Data type String

Default Null

JNDI name:

The JNDI name that is used to bind the connection factory into the application server's name space.

As a convention, use the fully qualified JNDI name; for example, in the form `jms/Name`, where *Name* is the logical name of the resource.

This name is used to link the platform binding information. The binding associates the resources defined by the deployment descriptor of the module to the actual (physical) resources bound into JNDI by the platform.

Data type String

Description:

A description of this connection factory for administrative purposes within IBM WebSphere Application Server.

Data type String
Default Null

Category:

A category used to classify or group this connection factory, for your IBM WebSphere Application Server administrative records.

Data type String

Component-managed Authentication Alias:

This alias specifies a user ID and password to be used to authenticate connection to a JMS provider for application-managed authentication.

This property provides a list of the J2C authentication data entry aliases that have been defined to WebSphere Application Server. You can select a data entry alias to be used to authenticate the creation of a new connection to the JMS provider.

If you have enabled global security for WebSphere Application Server, select the alias that specifies the user ID and password used to authenticate the creation of a new connection to the JMS provider. The use of this alias depends on the resource authentication (`res-auth`) setting declared in the connection factory resource reference of an application component's deployment descriptors.

Note: User IDs longer than 12 characters cannot be used for authentication with the embedded WebSphere JMS provider. For example, the default Windows NT user ID, **Administrator**, is not valid for use with embedded WebSphere messaging, because it contains 13 characters. Therefore, an authentication alias for a WebSphere JMS provider connection factory must specify a user ID no longer than 12 characters.

Data type Pick-list

Container-managed Authentication Alias:

This alias specifies a user ID and password to be used to authenticate connection to a JMS provider for container-managed authentication.

This property provides a list of the J2C authentication data entry aliases that have been defined to WebSphere Application Server. You can select a data entry alias to be used to authenticate the creation of a new connection to the JMS provider.

If you have enabled global security for WebSphere Application Server, select the alias that specifies the user ID and password used to authenticate the creation of a new connection to the JMS provider. The use of this alias depends on the resource authentication (res-auth) setting declared in the connection factory resource reference of an application component's deployment descriptors.

Note: User IDs longer than 12 characters cannot be used for authentication with the embedded WebSphere JMS provider. For example, the default Windows NT user ID, **Administrator**, is not valid for use with embedded WebSphere messaging, because it contains 13 characters. Therefore, an authentication alias for a WebSphere JMS provider connection factory must specify a user ID no longer than 12 characters.

Data type Pick-list

Mapping-Configuration Alias:

Allows users to select from the **Security > JAAS Configuration > Application Logins Configuration** list.

The **DefaultPrincipalMapping** JAAS configuration maps the authentication alias to the userid and password. You may define and use other mapping configurations. For more information about the mapping configurations, see "Java Authentication and Authorization service configuration entry settings" in the Information center.

Data type Pick-list

Node:

The WebSphere node name of the administrative node where the JMS server runs for this connection factory. Connections created by this factory connect to that JMS server.

Data type	String
Units	Enum
Default	Null
Range	Pull-down list of nodes in the WebSphere administrative domain.

XA Enabled:

Specifies whether the connection factory is for XA or non-XA coordination of messages and controls if the application server uses XA QCF/TCF. Enable XA if multiple resources are not used in the same transaction.

If you clear this checkbox property (for non-XA coordination), the JMS session is still enlisted in a transaction, but uses the resource manager local transaction calls (session.commit and session.rollback) instead of XA calls. This can lead to an improvement in performance. However, this means that only a single resource can be enlisted in a transaction in WebSphere Application Server.

5.1 + In WBI Server Foundation the last participant support enables you to enlist one non-XA resource with other XA-capable resources.

For a WebSphere Topic Connection Factory with the **Port** property set to DIRECT this property does not apply, and always adopts non-XA coordination.

Data type	Checkbox
Default	Selected (enabled for XA coordination)

Range**Selected**

The connection factory is enabled for XA-coordination of messages

Cleared

The connection factory is not enabled for XA coordination of messages

Recommended

Do not enable XA coordination when the message queue or topic received is the only resource in the transaction. Enable XA coordination when other resources, including other queues or topics, are involved.

Connection pool:

Specifies an optional set of connection pool settings.

Connection pool properties are common to all J2C connectors.

The application server pools connections and sessions with the JMS provider to improve performance. You need to configure the connection and session pool properties appropriately for your applications, otherwise you may not get the connection and session behavior that you want.

Change the size of the connection pool if concurrent server-side access to the JMS resource exceeds the default value. The size of the connection pool is set on a per queue or topic basis. See the following table for details.

Session pool:

An optional set of session pool settings.

This link provides a panel of optional connection pool properties, common to all J2C connectors.

The application server pools connections and sessions with the JMS provider to improve performance. You need to configure the connection and session pool properties appropriately for your applications, otherwise you may not get the connection and session behavior that you want.

Session pool settings:

Use this page to configure session pool settings.

This administrative console page is common to a range of resource types; for example, JMS queue connection factories. To view this page, the path depends on the type of resource, but generally you select an instance of the resource provider, then an instance of the resource type, then click **Session Pool**. For example: click **Resources > WebSphere JMS Providers > WebSphere Queue Connection Factories > connection_factory > Session Pool**.

Scope:

Specifies the level to which this resource definition is visible -- the cell, node, or server level.

Resources such as JDBC Providers, Namespace bindings, or shared libraries can be defined at multiple scopes, with resources defined at more specific scopes overriding duplicates which are defined at more general scopes.

Note that no matter what the scope of a defined resource, the resource's properties only apply at an individual server level. For example, if you define the scope of a data source at the Cell level, all users in that Cell can look up and use that data source, which is unique within that Cell. However, resource

property settings are local to each server in the Cell. For example, if you set **Max Connections** to 10, then each server in that Cell can have 10 connections.

Cell The most general scope. Resources defined at the Cell scope are visible from all Nodes and servers, unless they are overridden. To view resources defined in the cell scope, do not specify a server or a node name in the scope selection form.

Node The default scope for most resource types. Resources defined at the Node scope override any duplicates defined at the Cell scope and are visible to all servers on the same node, unless they are overridden at a server scope on that node. To view resources defined in a node scope, do not specify a server, but select a node name in the scope selection form.

Server The most specific scope for defining resources. Resources defined at the Server scope override any duplicate resource definitions defined at the Cell scope or parent Node scope and are visible only to a specific server. To view resources defined in a server scope, specify a server name as well as a node name in the scope selection form.

When resources are created, they are always created into the current scope selected in the panel. To view resources in other scopes, specify a different node or server in the scope selection form.

Data type String

Connection Timeout:

Specifies the interval, in seconds, after which a connection request times out and a `ConnectionWaitTimeoutException` is thrown.

The wait is necessary when the maximum value of connections (**Max Connections**) to a particular connection pool is reached. For example, if *Connection Timeout* is set to 300 and the maximum number of connections is reached, the Pool Manager waits for 300 seconds for an available physical connection. If a physical connection is *not* available within this time, the Pool Manager throws a `ConnectionWaitTimeoutException`. It usually does not make sense to retry the `getConnection()` method, because if a longer wait time is required, you should set the **Connection Timeout** setting to a higher value. Therefore, if this exception is caught by the application, the administrator should review the expected usage of the application and tune the connection pool and the database accordingly.

If *Connection Timeout* is set to 0, the Pool Manager waits as long as necessary until a connection is allocated (which happens when the number of connections falls below the value of **Max Connections**).

If *Max Connections* is set to 0, which enables an infinite number of physical connections, then the *Connection Timeout* value is ignored.

Data type	Integer
Units	Seconds
Default	180
Range	0 to max int

Max Connections:

Specifies the maximum number of physical connections that you can create in this pool.

These are the physical connections to the backend resource. Once this number is reached, no new physical connections are created and the requester waits until a physical connection that is currently in use returns to the pool, or a `ConnectionWaitTimeoutException` is thrown.

For example, if the Max Connections value is set to 5, and there are five physical connections in use, the pool manager waits for the amount of time specified in Connection Timeout for a physical connection to become free.

If Max Connections is set to 0, the Connection Timeout value is ignored.

For better performance, set the value for the connection pool lower than the value for the Max Connections option in the Web container. Lower settings, such as 10-30 connections, perform better than higher settings, such as 100.

If clones are used, one data pool exists for each clone. Knowing the number of data pools is important when configuring the database maximum connections.

You can use the Tivoli Performance Viewer to find the optimal number of connections in a pool. If the number of concurrent waiters is greater than 0, but the CPU load is not close to 100%, consider increasing the connection pool size. If the Percent Used value is consistently low under normal workload, consider decreasing the number of connections in the pool.

Data type	Integer
Default	10
Range	0 to max int

Min Connections:

Specifies the minimum number of physical connections to maintain.

Until this number is reached, the pool maintenance thread does not discard physical connections. However, no attempt is made to bring the number of connections up to this number. If you set a value for Aged Timeout, the minimum is not maintained. All connections with an expired age are discarded.

For example if the **Min Connections** value is set to 3, and one physical connection is created, the Unused Timeout thread does not discard that connection. By the same token, the thread does not automatically create two additional physical connections to reach the **Min Connections** setting.

Data type	Integer
Default	1
Range	0 to max int

Reap Time:

Specifies the interval, in seconds, between runs of the pool maintenance thread.

For example, if **Reap Time** is set to 60, the pool maintenance thread runs every 60 seconds. The Reap Time interval affects the accuracy of the **Unused Timeout** and **Aged Timeout** settings. The smaller the interval, the greater the accuracy. If the pool maintenance thread is enabled, set the Reap Time value less than the values of Unused Timeout and Aged Timeout. When the pool maintenance thread runs, it discards any connections remaining unused for longer than the time value specified in Unused Timeout, until it reaches the number of connections specified in **Min Connections**. The pool maintenance thread also discards any connections that remain active longer than the time value specified in Aged Timeout.

The Reap Time interval also affects performance. Smaller intervals mean that the pool maintenance thread runs more often and degrades performance.

To disable the pool maintenance thread set Reap Time to 0, or set both Unused Timeout and Aged Timeout to 0. The recommended way to disable the pool maintenance thread is to set Reap Time to 0, in

which case Unused Timeout and Aged Timeout are ignored. However, if Unused Timeout and Aged Timeout are set to 0, the pool maintenance thread runs, but only physical connections which timeout due to non-zero timeout values are discarded.

Data type	Integer
Units	Seconds
Default	180
Range	0 to max int

Unused Timeout:

Specifies the interval in seconds after which an unused or idle connection is discarded.

Set the Unused Timeout value higher than the Reap Timeout value for optimal performance. Unused physical connections are only discarded if the current number of connections not in use exceeds the **Min Connections** setting. For example, if the unused timeout value is set to 120, and the pool maintenance thread is enabled (Reap Time is not 0), any physical connection that remains unused for two minutes is discarded. Note that accuracy of this timeout, as well as performance, is affected by the **Reap Time** value. For more information, see Reap Time.

Data type	Integer
Units	Seconds
Default	1800
Range	0 to max int

Aged Timeout:

Specifies the interval in seconds before a physical connection is discarded.

Setting **Aged Timeout** to 0 supports active physical connections remaining in the pool indefinitely. Set the Aged Timeout value higher than the **Reap Timeout** value for optimal performance. For example, if the Aged Timeout value is set to 1200, and the Reap Time value is not 0, any physical connection that remains in existence for 1200 seconds (20 minutes) is discarded from the pool. Note that accuracy of this timeout, as well as performance, are affected by the Reap Time value. For more information, see Reap Time.

Data type	Integer
Units	Seconds
Default	0
Range	0 to max int

Purge Policy:

Specifies how to purge connections when a *stale connection* or *fatal connection error* is detected.

Valid values are **EntirePool** and **FailingConnectionOnly**. JCA data sources can have either option. WebSphere Version 4.0 data sources always have a purge policy of **EntirePool**.

Data type	String
Default	FailingConnectionOnly

Range

EntirePool

All connections in the pool are marked stale. Any connection not in use is immediately closed. A connection in use is closed and throws a *StaleConnectionException* during the next operation on that connection. Subsequent *getConnection* requests from the application result in new connections to the database opening. When using this purge policy, there is a slight possibility that some connections in the pool are closed unnecessarily when they are not stale. However, this is a rare occurrence. In most cases, a purge policy of EntirePool is the best choice.

FailingConnectionOnly

Only the connection that caused the *StaleConnectionException* is closed. Although this setting eliminates the possibility that valid connections are closed unnecessarily, it makes recovery from an application perspective more complicated. Because only the currently failing connection is closed, there is a good possibility that the next *getConnection* request from the application can return a connection from the pool that is also stale, resulting in more stale connection exceptions.

WebSphere topic connection factory collection:

The topic connection factories configured in the embedded WebSphere JMS provider for publish/subscribe messaging with JMS topics.

This panel shows a list of the WebSphere topic connection factories with a summary of their configuration properties.

To view this administrative console page, use the administrative console to complete the following steps:

1. In the navigation pane, expand **Resources-> WebSphere JMS Provider**.
2. In the content pane, select the JMS provider that you want to support the connection factory.
3. In the Additional Properties list in the contents pane, select **WebSphere Topic Connection Factory**.

To view or change the properties of a connection factory, select its name in the list displayed.

To act on one or more of the connection factories listed, click the check boxes next to the names of the objects that you want to act on, then use the buttons provided.

WebSphere topic connection factory settings:

Use this panel to view or change the configuration properties of the selected topic connection factory for use with the embedded WebSphere JMS provider. These configuration properties control how connections are created to the associated JMS topic destination.

A topic connection factory is used to create JMS connections to topic destinations. The topic connection factory is created by the associated JMS provider. A topic connection factory for the embedded WebSphere JMS provider has the following properties.

To view this administrative console page, click **Resources-> WebSphere JMS Provider-> (In content pane, under Additional Properties) WebSphere Topic Connection Factories-> connection_factory**

Scope:

Specifies the level to which this resource definition is visible -- the cell, node, or server level.

Resources such as JMS Providers, Namespace bindings, or shared libraries can be defined at multiple scopes, with resources defined at more specific scopes overriding duplicates which are defined at more general scopes.

Note that no matter what the scope of a defined resource, the resource's properties only apply at an individual server level. For example, if you define the scope of a data source at the Cell level, all users in that Cell can look up and use that data source, which is unique within that Cell. However, resource property settings are local to each server in the Cell. For example, if you set **Max Connections** to 10, then each server in that Cell can have 10 connections.

Cell The most general scope. Resources defined at the Cell scope are visible from all Nodes and servers, unless they are overridden. To view resources defined in the cell scope, do not specify a server or a node name in the scope selection form.

Node The default scope for most resource types. Resources defined at the Node scope override any duplicates defined at the Cell scope and are visible to all servers on the same node, unless they are overridden at a server scope on that node. To view resources defined in a node scope, do not specify a server, but select a node name in the scope selection form.

Server The most specific scope for defining resources. Resources defined at the Server scope override any duplicate resource definitions defined at the Cell scope or parent Node scope and are visible only to a specific server. To view resources defined in a server scope, specify a server name as well as a node name in the scope selection form.

When resources are created, they are always created into the current scope selected in the panel. To view resources in other scopes, specify a different node or server in the scope selection form.

Data type String

Name:

The name by which this queue connection factory is known for administrative purposes within IBM WebSphere Application Server. The name must be unique within the JMS connection factories across the WebSphere administrative domain.

Data type String
Default Null

JNDI name:

The JNDI name that is used to bind the topic connection factory into the application server's name space.

As a convention, use the fully qualified JNDI name; for example, in the form `jms/Name`, where *Name* is the logical name of the resource.

This name is used to link the platform binding information. The binding associates the resources defined by the deployment descriptor of the module to the actual (physical) resources bound into JNDI by the platform.

Data type String

Description:

A description of this topic connection factory for administrative purposes within IBM WebSphere Application Server.

Data type	String
Default	Null

Category:

A category used to classify or group this topic connection factory, for your IBM WebSphere Application Server administrative records.

Data type	String
------------------	--------

Node:

The WebSphere node name of the administrative node where the JMS server runs for this connection factory. Connections created by this factory connect to that JMS server.

Data type	Enum
Default	Null
Range	Pull-down list of nodes in the WebSphere administrative domain.

Port:

Which of the two ports that connections use to connect to the JMS Server. The QUEUED port is for full-function JMS publish/subscribe support, the DIRECT port is for non-persistent, non-transactional, non-durable subscriptions only.

Note: Message-driven beans cannot use the direct listener port for publish/subscribe support. Therefore, any topic connection factory configured with **Port** set to `Direct` cannot be used with message-driven beans.

Data type	Enum
Units	Not applicable
Default	QUEUED
Range	QUEUED The listener port used for full-function JMS-compliant, publish/subscribe support. DIRECT The listener port used for direct TCP/IP connection (non-transactional, non-persistent, and non-durable subscriptions only) for publish/subscribe support. The TCP/IP port numbers for these ports are defined on the WebSphere Internal JMS Server.

Component-managed Authentication Alias:

This alias specifies a user ID and password to be used to authenticate connection to a JMS provider for application-managed authentication.

This property provides a list of the J2C authentication data entry aliases that have been defined to WebSphere Application Server. You can select a data entry alias to be used to authenticate the creation of a new connection to the JMS provider.

If you have enabled global security for WebSphere Application Server, select the alias that specifies the user ID and password used to authenticate the creation of a new connection to the JMS provider. The use of this alias depends on the resource authentication (res-auth) setting declared in the connection factory resource reference of an application component's deployment descriptors.

Note: User IDs longer than 12 characters cannot be used for authentication with the embedded WebSphere JMS provider. For example, the default Windows NT user ID, **Administrator**, is not valid for use with embedded WebSphere messaging, because it contains 13 characters. Therefore, an authentication alias for a WebSphere JMS provider connection factory must specify a user ID no longer than 12 characters.

Container-managed Authentication Alias:

This alias specifies a user ID and password to be used to authenticate connection to a JMS provider for container-managed authentication.

This property provides a list of the J2C authentication data entry aliases that have been defined to WebSphere Application Server. You can select a data entry alias to be used to authenticate the creation of a new connection to the JMS provider.

If you have enabled global security for WebSphere Application Server, select the alias that specifies the user ID and password used to authenticate the creation of a new connection to the JMS provider. The use of this alias depends on the resource authentication (res-auth) setting declared in the connection factory resource reference of an application component's deployment descriptors.

Note: User IDs longer than 12 characters cannot be used for authentication with the embedded WebSphere JMS provider. For example, the default Windows NT user ID, **Administrator**, is not valid for use with embedded WebSphere messaging, because it contains 13 characters. Therefore, an authentication alias for a WebSphere JMS provider connection factory must specify a user ID no longer than 12 characters.

Mapping-Configuration Alias:

Allows users to select from the **Security > JAAS Configuration > Application Logins Configuration** list.

The **DefaultPrincipalMapping** JAAS configuration maps the authentication alias to the userid and password. You may define and use other mapping configurations. For more information about the mapping configurations, see "Java Authentication and Authorization service configuration entry settings" in the Information center.

Data type Pick-list

Clone Support:

Select this checkbox to enable clone support to allow the same durable subscription across topic clones.

Data type Enum
Default Cleared

Range**Selected**

Clone support is enabled.

Cleared

Clone support is disabled.

If you select this property, you must also specify a value for the **Client ID** property.

Client ID:

The JMS client identifier used for connections to the queue manager.

Data type

String

Range

A valid JMS client ID

XA Enabled:

Specifies whether the connection factory is for XA or non-XA coordination of messages and controls if the application server uses XA QCF/TCF. Enable XA if multiple resources are not used in the same transaction.

If you clear this checkbox property (for non-XA coordination), the JMS session is still enlisted in a transaction, but uses the resource manager local transaction calls (`session.commit` and `session.rollback`) instead of XA calls. This can lead to an improvement in performance. However, this means that only a single resource can be enlisted in a transaction in WebSphere Application Server.

5.1+ In WBI Server Foundation the last participant support enables you to enlist one non-XA resource with other XA-capable resources.

For a WebSphere Topic Connection Factory with the **Port** property set to DIRECT this property does not apply, and always adopts non-XA coordination.

Data type

Checkbox

Default

Selected (enabled for XA coordination)

Range**Selected**

The connection factory is enabled for XA-coordination of messages

Cleared

The connection factory is not enabled for XA coordination of messages

Recommended

Do not enable XA coordination when the message queue or topic received is the only resource in the transaction. Enable XA coordination when other resources, including other queues or topics, are involved.

Connection pool:

Specifies an optional set of connection pool settings.

Connection pool properties are common to all J2C connectors.

The application server pools connections and sessions with the JMS provider to improve performance. You need to configure the connection and session pool properties appropriately for your applications, otherwise you may not get the connection and session behavior that you want.

Change the size of the connection pool if concurrent server-side access to the JMS resource exceeds the default value. The size of the connection pool is set on a per queue or topic basis. See the following table for details.

Session pool:

An optional set of session pool settings.

This link provides a panel of optional connection pool properties, common to all J2C connectors.

The application server pools connections and sessions with the JMS provider to improve performance. You need to configure the connection and session pool properties appropriately for your applications, otherwise you may not get the connection and session behavior that you want.

WebSphere Queue destination collection:

The queue destinations configured in the embedded WebSphere JMS provider for point-to-point messaging with JMS queues.

This panel shows a list of the WebSphere queue destinations with a summary of their configuration properties.

To view this administrative console page, use the administrative console to complete the following steps:

1. In the navigation pane, expand **Resources-> WebSphere JMS Provider**.
2. In the content pane, select the JMS provider that you want to support the queue destination.
3. In the Additional Properties list in the contents pane, select **WebSphere Queue Destination**.

To view or change the properties of a queue destination, select its name in the list displayed.

To act on one or more of the queue destinations listed, click the check boxes next to the names of the objects that you want to act on, then use the buttons provided.

WebSphere queue settings:

Use this panel to view or change the configuration properties of the selected queue destination for use with the WebSphere JMS provider.

A queue destination is used to configure the properties of a JMS queue. Connections to the queue are created by the associated queue connection factory for the embedded WebSphere JMS provider. A queue for use with the internal WebSphere JMS provider has the following properties.

To view this administrative console page, click **Resources-> WebSphere JMS Provider-> (In content pane, under Additional Properties) WebSphere Queue Destinations-> *destination_name***

Scope:

Specifies the level to which this resource definition is visible -- the cell, node, or server level.

Resources such as JMS Providers, Namespace bindings, or shared libraries can be defined at multiple scopes, with resources defined at more specific scopes overriding duplicates which are defined at more general scopes.

Note that no matter what the scope of a defined resource, the resource's properties only apply at an individual server level. For example, if you define the scope of a data source at the Cell level, all users in that Cell can look up and use that data source, which is unique within that Cell. However, resource

property settings are local to each server in the Cell. For example, if you set **Max Connections** to 10, then each server in that Cell can have 10 connections.

Cell The most general scope. Resources defined at the Cell scope are visible from all Nodes and servers, unless they are overridden. To view resources defined in the cell scope, do not specify a server or a node name in the scope selection form.

Node The default scope for most resource types. Resources defined at the Node scope override any duplicates defined at the Cell scope and are visible to all servers on the same node, unless they are overridden at a server scope on that node. To view resources defined in a node scope, do not specify a server, but select a node name in the scope selection form.

Server The most specific scope for defining resources. Resources defined at the Server scope override any duplicate resource definitions defined at the Cell scope or parent Node scope and are visible only to a specific server. To view resources defined in a server scope, specify a server name as well as a node name in the scope selection form.

When resources are created, they are always created into the current scope selected in the panel. To view resources in other scopes, specify a different node or server in the scope selection form.

Data type String

Name:

The name by which the queue is known for administrative purposes within IBM WebSphere Application Server.

To enable applications to use this queue, you must add the queue name to the list of queue names in the configuration of the JMS servers that host the queue.

Data type String

JNDI name:

The JNDI name that is used to bind the queue into the application server's name space.

As a convention, use the fully qualified JNDI name; for example, in the form `jms/Name`, where *Name* is the logical name of the resource.

This name is used to link the platform binding information. The binding associates the resources defined by the deployment descriptor of the module to the actual (physical) resources bound into JNDI by the platform.

Data type String

Description:

A description of the queue, for administrative purposes

Data type String

Default Null

Category:

A category used to classify or group this queue, for your IBM WebSphere Application Server administrative records.

Data type String

Persistence:

Whether all messages sent to the destination are persistent, non-persistent, or have their persistence defined by the application

Data type	Enum
Default	APPLICATION DEFINED
Range	APPLICATION DEFINED Messages on the destination have their persistence defined by the application that put them onto the queue. NON-PERSISTENT Messages on the destination are not persistent. PERSISTENT Messages on the destination are persistent. When a persistent message is put to a queue, all of the message data is written to the messaging log (under the <i>embedded_messaging_install</i> log directory) to make recovery of the message possible.

Priority:

Whether the message priority for this destination is defined by the application or the **Specified priority** property

Data type	Enum
Units	Not applicable
Default	APPLICATION DEFINED
Range	APPLICATION DEFINED The priority of messages on this destination is defined by the application that put them onto the destination. QUEUE DEFINED [WebSphere MQ destination only] Messages on the destination have their persistence defined by the WebSphere MQ queue definition properties. SPECIFIED The priority of messages on this destination is defined by the Specified priority property. <i>If you select this option, you must define a priority on the Specified priority property.</i>

Specified priority:

If the **Priority** property is set to Specified, type here the message priority for this queue, in the range 0 (lowest) through 9 (highest)

If the **Priority** property is set to Specified, messages sent to this queue have the priority value specified by this property.

Data type	Integer
Units	Message priority level
Default	0

Range 0 (lowest priority) through 9 (highest priority)

Expiry:

Whether the expiry timeout for this queue is defined by the application or the **Specified expiry** property, or messages on the queue never expire (have an unlimited expiry timeout)

Data type	Enum
Units	Not applicable
Default	APPLICATION DEFINED
Range	APPLICATION DEFINED The expiry timeout for messages on this queue is defined by the application that put them onto the queue. UNLIMITED Messages on this queue have no expiry timeout, so those messages never expire. SPECIFIED The expiry timeout for messages on this queue is defined by the Specified expiry property. <i>If you select this option, you must define a timeout on the Specified expiry property.</i>

Specified expiry:

If the **Expiry timeout** property is set to *Specified*, type here the number of milliseconds (greater than 0) after which messages on this queue expire

Data type	Integer
Units	Milliseconds
Default	0
Range	Greater than or equal to 0 <ul style="list-style-type: none">• 0 indicates that messages never timeout• Other values are an integer number of milliseconds

WebSphere topic destination collection:

The topic destinations configured in the embedded WebSphere JMS provider for publish/subscribe messaging with JMS topics. Use this panel to create or delete topic destinations, or to select a topic destination to view or change its configuration properties.

This panel shows a list of the WebSphere topic destinations with a summary of their configuration properties.

To view this administrative console page, use the administrative console to complete the following steps:

1. In the navigation pane, expand **Resources-> WebSphere JMS Provider**.
2. In the content pane, select the JMS provider that you want to support the topic destination.
3. In the Additional Properties list in the contents pane, select **WebSphere Topic Destination**.

To view or change the properties of a topic destination, select its name in the list displayed.

To act on one or more of the topic destinations listed, click the check boxes next to the names of the objects that you want to act on, then use the buttons provided.

WebSphere topic settings:

Use this panel to view or change the configuration properties of the selected topic destination for use with the embedded WebSphere JMS provider.

A topic destination is used to configure the properties of a JMS topic for the associated JMS provider. Connections to the topic are created by the associated topic connection factory. A topic for use with the embedded WebSphere JMS provider has the following properties.

To view this administrative console page, click **Resources-> WebSphere JMS Provider-> (In content pane, under Additional Properties) WebSphere Topic destinations-> *destination_name***

Scope:

Specifies the level to which this resource definition is visible -- the cell, node, or server level.

Resources such as JMS Providers, Namespace bindings, or shared libraries can be defined at multiple scopes, with resources defined at more specific scopes overriding duplicates which are defined at more general scopes.

Note that no matter what the scope of a defined resource, the resource's properties only apply at an individual server level. For example, if you define the scope of a data source at the Cell level, all users in that Cell can look up and use that data source, which is unique within that Cell. However, resource property settings are local to each server in the Cell. For example, if you set **Max Connections** to 10, then each server in that Cell can have 10 connections.

Cell The most general scope. Resources defined at the Cell scope are visible from all Nodes and servers, unless they are overridden. To view resources defined in the cell scope, do not specify a server or a node name in the scope selection form.

Node The default scope for most resource types. Resources defined at the Node scope override any duplicates defined at the Cell scope and are visible to all servers on the same node, unless they are overridden at a server scope on that node. To view resources defined in a node scope, do not specify a server, but select a node name in the scope selection form.

Server The most specific scope for defining resources. Resources defined at the Server scope override any duplicate resource definitions defined at the Cell scope or parent Node scope and are visible only to a specific server. To view resources defined in a server scope, specify a server name as well as a node name in the scope selection form.

When resources are created, they are always created into the current scope selected in the panel. To view resources in other scopes, specify a different node or server in the scope selection form.

Data type String

Name:

The name by which the topic is known for administrative purposes within IBM WebSphere Application Server.

Data type String

JNDI name:

The JNDI name that is used to bind the topic into the application server's name space.

As a convention, use the fully qualified JNDI name; for example, in the form *jms/Name*, where *Name* is the logical name of the resource.

This name is used to link the platform binding information. The binding associates the resources defined by the deployment descriptor of the module to the actual (physical) resources bound into JNDI by the platform.

Data type String

Description:

A description of the topic, for administrative purposes within IBM WebSphere Application Server.

Data type String
Default Null

Category:

A category used to classify or group this topic, for your IBM WebSphere Application Server administrative records.

Data type String

Topic:

The name of the topic as defined to the JMS provider.

Data type String
Default Null
Range The topic value can be dot notation and include wildcard characters.

Persistence:

Whether all messages sent to the destination are persistent, non-persistent, or have their persistence defined by the application

Data type Enum
Default APPLICATION DEFINED
Range **APPLICATION DEFINED**
Messages on the destination have their persistence defined by the application that put them onto the queue.
NON-PERSISTENT
Messages on the destination are not persistent.
PERSISTENT
Messages on the destination are persistent.
QUEUE DEFINED
[WebSphere MQ destination only] Messages on the destination have their persistence defined by the WebSphere MQ queue definition properties.

Priority:

Whether the message priority for this destination is defined by the application or the **Specified priority** property

Data type
Units
Default
Range

Enum
Not applicable
APPLICATION DEFINED
APPLICATION DEFINED
The priority of messages on this destination is defined by the application that put them onto the destination.
QUEUE DEFINED
[WebSphere MQ destination only] Messages on the destination have their persistence defined by the WebSphere MQ queue definition properties.
SPECIFIED
The priority of messages on this destination is defined by the **Specified priority** property. *If you select this option, you must define a priority on the **Specified priority** property.*

Specified priority:

If the **Priority** property is set to Specified, type here the message priority for this queue, in the range 0 (lowest) through 9 (highest)

If the **Priority** property is set to Specified, messages sent to this queue have the priority value specified by this property.

Data type
Units
Default
Range

Integer
Message priority level
0
0 (lowest priority) through 9 (highest priority)

Expiry:

Whether the expiry timeout for this queue is defined by the application or the **Specified expiry** property, or messages on the queue never expire (have an unlimited expiry timeout)

Data type
Units
Default
Range

Enum
Not applicable
APPLICATION DEFINED
APPLICATION DEFINED
The expiry timeout for messages on this queue is defined by the application that put them onto the queue.
UNLIMITED
Messages on this queue have no expiry timeout, so those messages never expire.
SPECIFIED
The expiry timeout for messages on this queue is defined by the **Specified expiry** property. *If you select this option, you must define a timeout on the **Specified expiry** property.*

Specified expiry:

If the **Expiry timeout** property is set to Specified, type here the number of milliseconds (greater than 0) after which messages on this queue expire

Data type	Integer
Units	Milliseconds
Default	0
Range	Greater than or equal to 0 <ul style="list-style-type: none"> • 0 indicates that messages never timeout • Other values are an integer number of milliseconds

WebSphere MQ queue connection factory collection:

The queue connection factories configured in the WebSphere MQ JMS provider for point-to-point messaging with JMS queues.

This panel shows a list of the WebSphere MQ queue connection factories with a summary of their configuration properties.

To view this administrative console page, use the administrative console to complete the following steps:

1. In the navigation pane, expand **Resources-> WebSphere MQ JMS Provider**.
2. In the content pane, select the JMS provider that you want to support the connection factory.
3. In the Additional Properties list in the contents pane, select **WebSphere MQ Queue Connection Factory**.

To view or change the properties of a connection factory, select its name in the list displayed.

To act on one or more of the connection factories listed, click the check boxes next to the names of the objects that you want to act on, then use the buttons provided.

WebSphere MQ queue connection factory settings:

Use this panel to view or change the configuration properties of the selected queue connection factory for use with the WebSphere MQ JMS provider. These configuration properties control how connections are created to the associated JMS queue destination.

A queue connection factory is used to create JMS connections to queue destinations. The queue connection factory is created by the WebSphere MQ JMS provider. A queue connection factory for the WebSphere MQ JMS provider has the following properties.

Note:

- The property values that you specify must match the values that you specified when configuring WebSphere MQ for JMS resources. For more information about configuring WebSphere MQ JMS resources, see the *WebSphere MQ Using Java* book, and the *WebSphere MQ System Administration* book, SC33-1873, which are available from the WebSphere MQ messaging platform-specific books Web page at <http://www-3.ibm.com/software/ts/mqseries/library/manualsa/manuals/platspecific.html>, the IBM Publications Center, or from the WebSphere MQ collection kit, SK2T-0730.
- In WebSphere MQ, names can have a maximum of 48 characters, with the exception of channels which have a maximum of 20 characters.
- **5.1+** You can use the Custom properties page to define WebSphere MQ Secure Sockets Layer (SSL) properties. For more information about setting those properties, see Custom properties.

To view this administrative console page, click **Resources-> WebSphere MQ JMS Provider-> (In content pane, under Additional Properties) WebSphere MQ Queue Connection Factories-> connection_factory**

Name:

The name by which this queue connection factory is known for administrative purposes within IBM WebSphere Application Server. The name must be unique within the JMS connection factories across the WebSphere administrative domain.

Data type String

JNDI name:

The JNDI name that is used to bind the connection factory into the application server's name space.

As a convention, use the fully qualified JNDI name; for example, in the form `jms/Name`, where *Name* is the logical name of the resource.

This name is used to link the platform binding information. The binding associates the resources defined by the deployment descriptor of the module to the actual (physical) resources bound into JNDI by the platform.

Data type String

Description:

A description of this connection factory for administrative purposes within IBM WebSphere Application Server.

Data type String
Default Null

Category:

A category used to classify or group this connection factory, for your IBM WebSphere Application Server administrative records.

Data type String

Component-managed Authentication Alias:

This alias specifies a user ID and password to be used to authenticate connection to a JMS provider for application-managed authentication.

This property provides a list of the J2C authentication data entry aliases that have been defined to WebSphere Application Server. You can select a data entry alias to be used to authenticate the creation of a new connection to the JMS provider.

If you have enabled global security for WebSphere Application Server, select the alias that specifies the user ID and password used to authenticate the creation of a new connection to the JMS provider. The use of this alias depends on the resource authentication (`res-auth`) setting declared in the connection factory resource reference of an application component's deployment descriptors.

Restriction:

1. User IDs longer than 12 characters cannot be used for authentication with the embedded WebSphere JMS provider. For example, the default Windows NT user ID, **Administrator**, is not valid for use with embedded WebSphere messaging, because it contains 13 characters. Therefore, an authentication alias for a WebSphere JMS provider connection factory must specify a user ID no longer than 12 characters.

2. If you want to use a WebSphere MQ JMS Provider JMS connection when using Bindings transport mode, you set the property **Transport type**=BINDINGS on the WebSphere MQ Queue Connection Factory. You must also choose one of the following options:
 - To use security credentials, ensure that the user specified is the currently logged on user for the WebSphere Application Server process. If the user specified is not the current logged on user for the WebSphere Application Server process, then the WebSphere MQ JMS Bindings authentication throws the error MQJMS2013 invalid security authentication supplied for MQQueueManager.
 - Do not specify security credentials. On the WebSphere MQ Connection Factory, ensure that both the **Component-managed Authentication Alias** and the **Container-managed Authentication Alias** properties are not set.

Container-managed Authentication Alias:

This alias specifies a user ID and password to be used to authenticate connection to a JMS provider for container-managed authentication.

This property provides a list of the J2C authentication data entry aliases that have been defined to WebSphere Application Server. You can select a data entry alias to be used to authenticate the creation of a new connection to the JMS provider.

If you have enabled global security for WebSphere Application Server, select the alias that specifies the user ID and password used to authenticate the creation of a new connection to the JMS provider. The use of this alias depends on the resource authentication (res-auth) setting declared in the connection factory resource reference of an application component's deployment descriptors.

Restriction:

1. User IDs longer than 12 characters cannot be used for authentication with the embedded WebSphere JMS provider. For example, the default Windows NT user ID, **Administrator**, is not valid for use with embedded WebSphere messaging, because it contains 13 characters. Therefore, an authentication alias for a WebSphere JMS provider connection factory must specify a user ID no longer than 12 characters.
2. If you want to use a WebSphere MQ JMS Provider JMS connection when using Bindings transport mode, you set the property **Transport type**=BINDINGS on the WebSphere MQ Queue Connection Factory. You must also choose one of the following options:
 - To use security credentials, ensure that the user specified is the currently logged on user for the WebSphere Application Server process. If the user specified is not the current logged on user for the WebSphere Application Server process, then the WebSphere MQ JMS Bindings authentication throws the error MQJMS2013 invalid security authentication supplied for MQQueueManager.
 - Do not specify security credentials. On the WebSphere MQ Connection Factory, ensure that both the **Component-managed Authentication Alias** and the **Container-managed Authentication Alias** properties are not set.

Queue manager:

The name of the WebSphere MQ queue manager for this connection factory. Connections created by this factory connect to that queue manager.

Data type	String
Default	Null
Range	A valid WebSphere MQ queue manager name, as 1 through 48 ASCII characters

Host:

The name of the host on which the WebSphere MQ queue manager runs, for client connection only.

Data type	String
Default	Null
Range	A valid TCP/IP hostname

Port:

The TCP/IP port number used for connection to the WebSphere MQ queue manager, for client connection only.

This port must be configured on the WebSphere MQ queue manager.

Data type	Integer
Default	Null
Range	A valid TCP/IP port number, configured on the WebSphere MQ queue manager.

Channel:

The name of the channel used for connection to the WebSphere MQ queue manager, for client connection only.

Data type	String
Default	Null
Range	1 through 20 ASCII characters

Transport type:

Specifies whether the WebSphere MQ client connection or JNI bindings are used for connection to the WebSphere MQ queue manager. The external JMS provider controls the communication protocols between JMS clients and JMS servers. Tune the transport type when you are using non-ASF non-persistent, non-durable, non-transactional messaging or when you want to satisfy security issues and the client is local to the queue manager node.

Data type	Enum
Units	Not applicable
Default	BINDINGS
Range	BINDINGS JNI bindings are used to connect to the queue manager. BINDINGS is a shared memory protocol and can only be used when the queue manager is on the same node as the JMS client and comes at some security risks that should be addressed through the use of EJB roles.
	CLIENT WebSphere MQ client connection is used to connect to the queue manager. CLIENT is a typical TCP-based protocol.
Recommended	BINDINGS is faster by 30% or more, but it lacks security. When you have security concerns, BINDINGS is more desirable than CLIENT.

Model queue definition:

The name of the model queue definition that can be used by the queue manager to create temporary queues if a queue requested does not already exist.

Data type	String
Default	Null
Range	1 through 48 ASCII characters

Client ID:

The JMS client identifier used for connections to the WebSphere MQ queue manager.

Data type	String
Default	Null

CCSID:

The coded character set identifier for use with the WebSphere MQ queue manager.

This coded character set identifier (CCSID) must be one of the CCSIDs supported by WebSphere MQ.

Data type	String
Units	Integer
Default	Null
Range	1 through 65535

For more information about supported CCSIDs, and about converting between message data from one coded character set to another, see the *WebSphere MQ System Administration* and the *WebSphere MQ Application Programming Reference* books. These are available from the WebSphere MQ messaging multiplatform and platform-specific books Web pages; for example, at <http://www-3.ibm.com/software/ts/mqseries/library/manualsa/manuals/platspecific.html>, the IBM Publications Center, or from the WebSphere MQ collection kit, SK2T-0730.

Message retention:

Select this check box to specify that unwanted messages are to be left on the queue. Otherwise, unwanted messages are dealt with according to their disposition options.

Data type	Enum
Units	Not applicable
Default	Cleared
Range	Selected Unwanted messages are left on the queue. Cleared Unwanted messages are dealt with according to their disposition options.

XA Enabled:

Specifies whether the connection factory is for XA or non-XA coordination of messages and controls if the application server uses XA QCF/TCF. Enable XA if multiple resources are not used in the same transaction.

If you set this property to `NON_XA`, the JMS session is still enlisted in a transaction, but uses the resource manager local transaction calls (`session.commit` and `session.rollback`) instead of XA calls. This can lead to an improvement in performance. However, this means that only a single resource can be enlisted in a transaction in WebSphere Application Server.

5.1 + In WBI Server Foundation the last participant support enables you to enlist one non-XA resource with other XA-capable resources.

Data type	Enum
Units	Not applicable
Default	XA enabled
Range	<p>XA The connection factory is for XA-coordination of messages</p> <p>NON_XA The connection factory is for non-XA coordination of messages</p>
Recommended	Do not enable XA when the message queue received is the only resource in the transaction. Enable XA when other resources, including other queues or topics, are involved.

Connection pool:

Specifies an optional set of connection pool settings.

Connection pool properties are common to all J2C connectors.

The application server pools connections and sessions with the JMS provider to improve performance. This is independent from any WebSphere MQ connection pooling. You need to configure the connection and session pool properties appropriately for your applications, otherwise you may not get the connection and session behavior that you want.

Change the size of the connection pool if concurrent server-side access to the JMS resource exceeds the default value. The size of the connection pool is set on a per queue or topic basis.

Session pool:

An optional set of session pool settings.

This link provides a panel of optional connection pool properties, common to all J2C connectors.

The application server pools connections and sessions with the JMS provider to improve performance. This is independent from any WebSphere MQ connection pooling. You need to configure the connection and session pool properties appropriately for your applications, otherwise you may not get the connection and session behavior that you want.

Custom properties:

An optional set of name and value pairs for custom properties passed to WebSphere MQ.

You can use the Custom properties page to define the following WebSphere MQ Secure Sockets Layer (SSL) properties. These properties apply only if you set the property **Transport type=CLIENT**.

- SSLCIPHERSUITE
- SSLCRL
- SSLPEERNAME

For more information about setting these properties, see the section SSL properties in the *WebSphere MQ Using Java* book.

WebSphere MQ queue connection factory custom properties:

Use this panel to view or change an optional set of name and value pairs for custom properties of the selected queue connection factory for use with the WebSphere MQ JMS provider. These custom properties are passed to WebSphere MQ.

To view this administrative console page, click **Resources-> WebSphere MQ JMS Provider-> (In content pane, under Additional Properties) WebSphere MQ Queue Connection Factories-> connection_factory-> Custom Properties**

You can use the Custom properties page to define the following WebSphere MQ Secure Sockets Layer (SSL) properties. These properties apply only if you set the property **Transport type=CLIENT** on the connection factory.

- SSLCIPHERSUITE
- SSLCRL
- SSLPEERNAME

For more information about setting these properties for WebSphere MQ, see the section SSL properties in the *WebSphere MQ Using Java* book.

SSLCIPHERSUITE:

The cipher suite to use for SSL connection.

Set this property to a valid cipher suite provided by your JSSE provider; it must match the CipherSpec named on the SVRCONN channel named by the **Channel** property.

You must set this property if the SSLPEERNAME is to be set.

SSLCRL:

A list of zero or more CRL (Certificate Revocation List) servers used to check for SSL certificate revocation. (Use of this property requires a WebSphere MQ JVM at Java 2 version 1.4.)

The value is a space-delimited list of entries of the form:

`ldap://hostname:[port]`

optionally followed by a single / (forward slash). If *port* is omitted, the default LDAP port of 389 is assumed. At connect-time, the SSL certificate presented by the server is checked against the specified CRL servers. For more information about CRL security, see the section “Working with Certificate Revocation Lists” in the *WebSphere MQ Security book*; for example at:

<http://publibfp.boulder.ibm.com/epubs/html/csqzas01/csqzas012w.htm#IDX2254>.

SSLPEERNAME:

For SSL, a *distinguished name* skeleton that must match the name provided by the WebSphere MQ queue manager. The distinguished name is used to check the identifying certificate presented by the server at connect-time.

If SSLPEERNAME is not set, such checking is performed. SSLPEERNAME is ignored if SSLCIPHERSUITE is not specified.

The SSLPEERNAME property is a list of attribute name and value pairs separated by commas or semicolons. For example:

```
SSLPEERNAME(CN=QMGR.*, OU=IBM, OU=WEBSHERE)
```

The example given checks the identifying certificate presented by the server at connect-time. For the connection to succeed, the certificate must have a Common Name beginning QMGR., and must have at least two Organizational Unit names, the first of which is IBM and the second WEBSHERE. Checking is not case-sensitive.

For more details about distinguished names and their use with WebSphere MQ, see the *WebSphere MQ Security* book; for example, the section “Distinguished Names” at <http://publibfp.boulder.ibm.com/epubs/html/csqzas01/csqzas010p.htm#HDRDCDN>.

WebSphere MQ topic connection factory collection:

The topic connection factories configured in the WebSphere MQ JMS provider for publish/subscribe messaging with JMS topics.

This panel shows a list of the WebSphere MQ topic connection factories with a summary of their configuration properties.

To view this administrative console page, use the administrative console to complete the following steps:

1. In the navigation pane, expand **Resources-> WebSphere MQ JMS Provider**.
2. In the content pane, select the JMS provider that you want to support the connection factory.
3. In the Additional Properties list in the contents pane, select **WebSphere MQ Topic Connection Factory**.

To view or change the properties of a connection factory, select its name in the list displayed.

To act on one or more of the connection factories listed, click the check boxes next to the names of the objects that you want to act on, then use the buttons provided.

WebSphere MQ topic connection factory settings:

Use this panel to view or change the configuration properties of the selected topic connection factory for use with the WebSphere MQ JMS provider. These configuration properties control how connections are created to the associated JMS topic destination.

A topic connection factory is used to create JMS connections to topic destinations. The topic connection factory is created by the WebSphere MQ JMS provider. A topic connection factory for the WebSphere MQ JMS provider has the following properties.

Note:

- The property values that you specify must match the values that you specified when configuring WebSphere MQ JMS resources. For more information about configuring WebSphere MQ JMS resources, see the *WebSphere MQ Using Java* book.
- In WebSphere MQ, names can have a maximum of 48 characters, with the exception of channels which have a maximum of 20 characters.
- **5.1+** You can use the Custom properties page to define a subset of WebSphere MQ properties. For more information about setting those properties, see Custom properties.

To view this administrative console page, click **Resources-> WebSphere MQ JMS Provider-> (In content pane, under Additional Properties) WebSphere MQ Topic Connection Factories-> connection_factory**

Name:

The name by which this topic connection factory is known for administrative purposes within IBM WebSphere Application Server. The name must be unique within the JMS provider.

Data type String

JNDI name:

The JNDI name that is used to bind the topic connection factory into the application server's name space.

As a convention, use the fully qualified JNDI name; for example, in the form `jms/Name`, where *Name* is the logical name of the resource.

This name is used to link the platform binding information. The binding associates the resources defined by the deployment descriptor of the module to the actual (physical) resources bound into JNDI by the platform.

Data type String

Description:

A description of this topic connection factory for administrative purposes within IBM WebSphere Application Server.

Data type String
Default Null

Category:

A category used to classify or group this topic connection factory, for your IBM WebSphere Application Server administrative records.

Data type String

Component-managed Authentication Alias:

This alias specifies a user ID and password to be used to authenticate connection to a JMS provider for application-managed authentication.

This property provides a list of the J2C authentication data entry aliases that have been defined to WebSphere Application Server. You can select a data entry alias to be used to authenticate the creation of a new connection to the JMS provider.

If you have enabled global security for WebSphere Application Server, select the alias that specifies the user ID and password used to authenticate the creation of a new connection to the JMS provider. The use of this alias depends on the resource authentication (`res-auth`) setting declared in the connection factory resource reference of an application component's deployment descriptors.

Note: User IDs longer than 12 characters cannot be used for authentication with the embedded WebSphere JMS provider. For example, the default Windows NT user ID, **Administrator**, is not valid for use with embedded WebSphere messaging, because it contains 13 characters. Therefore, an authentication alias for a WebSphere JMS provider connection factory must specify a user ID no longer than 12 characters.

Container-managed Authentication Alias:

This alias specifies a user ID and password to be used to authenticate connection to a JMS provider for container-managed authentication.

This property provides a list of the J2C authentication data entry aliases that have been defined to WebSphere Application Server. You can select a data entry alias to be used to authenticate the creation of a new connection to the JMS provider.

If you have enabled global security for WebSphere Application Server, select the alias that specifies the user ID and password used to authenticate the creation of a new connection to the JMS provider. The use of this alias depends on the resource authentication (res-auth) setting declared in the connection factory resource reference of an application component's deployment descriptors.

Note: User IDs longer than 12 characters cannot be used for authentication with the embedded WebSphere JMS provider. For example, the default Windows NT user ID, **Administrator**, is not valid for use with embedded WebSphere messaging, because it contains 13 characters. Therefore, an authentication alias for a WebSphere JMS provider connection factory must specify a user ID no longer than 12 characters.

Queue manager:

The name of the WebSphere MQ queue manager for this connection factory. Connections created by this factory connect to that queue manager.

Data type	String
Default	Null
Range	A valid WebSphere MQ queue manager name, as 1 through 48 ASCII characters

Host:

The name of the host on which the WebSphere MQ queue manager runs, for client connection only.

Data type	String
Default	Null
Range	A valid TCP/IP hostname

Port:

The TCP/IP port number used for connection to the WebSphere MQ queue manager, for client connection only.

This port must be configured on the WebSphere MQ queue manager.

Data type	Integer
Default	Null
Range	A valid TCP/IP port number, configured on the WebSphere MQ queue manager.

Channel:

The name of the channel used for connection to the WebSphere MQ queue manager, for client connection only.

Data type	String
Default	Null

Range 1 through 20 ASCII characters

Transport type:

Specifies whether the WebSphere MQ client connection or JNI bindings are used for connection to the WebSphere MQ queue manager. The external JMS provider controls the communication protocols between JMS clients and JMS servers. Tune the transport type when you are using non-ASF nonpersistent, non-durable, non-transactional messaging or when you want to satisfy security issues and the client is local to the queue manager node.

Data type	Enum
Units	Not applicable
Default	BINDINGS
Range	BINDINGS JNI bindings are used to connect to the queue manager. BINDINGS is a shared memory protocol and can only be used when the queue manager is on the same node as the JMS client and comes at some security risks that should be addressed through the use of EJB roles. CLIENT WebSphere MQ client connection is used to connect to the queue manager. CLIENT is a typical TCP-based protocol. DIRECT For WebSphere MQ Event Broker using DIRECT mode. DIRECT is a lightweight sockets protocol used in non-transactional, non-durable and non-persistent Publish/Subscribe messaging. DIRECT works only for clients and message-driven beans using the non-ASF protocol.

Recommended DIRECT is the fastest transport type and should be used where possible. Use BINDINGS when you want to satisfy additional security tasks and the queue manager is local to the JMS client. QUEUED is fallback for all other cases. **Note:** WebSphere MQ 5.3 before CSD2 with the DIRECT setting can lose messages when used with message-driven beans and under load. This also happens with client-side based applications unless the broker's maxClientQueueSize is set to 0. You can set this to 0 with the command `#wempschangeproperties WAS_nodeName_server1 -e default -o DynamicSubscriptionEngine -n maxClientQueueSize -v 0 -x executionGroupUUID`, where executionGroupUUID can be found by starting the broker and looking in the Event Log/Applications for event 2201. This value is usually ffffffff-0000-0000-000000000000.

Broker control queue:

The name of the broker's control queue, to which all command messages (except publications and requests to delete publications) are sent

The name of the broker's control queue. Publisher and subscriber applications, and other brokers, send all command messages (except publications and requests to delete publications) to this queue.

Data type	String
Units	En_US ASCII characters
Default	Null
Range	1 through 48 ASCII characters

Broker queue manager:

The name of the WebSphere MQ queue manager that provides the publish/subscribe message broker.

Data type	String
Units	En_US ASCII characters
Default	Null
Range	1 through 48 ASCII characters

Broker publication queue:

The name of the broker's input queue that receives all publication messages for the default stream

The name of the broker's input queue (stream queue) that receives all publication messages for the default stream. Applications can also send requests to delete publications on the default stream to this queue.

Data type	String
Units	En_US ASCII characters
Default	Null
Range	1 through 48 ASCII characters

Broker subscription queue:

The name of the broker's queue from which non-durable subscription messages are retrieved

The name of the broker's queue from which non-durable subscription messages are retrieved. The subscriber specifies the name of the queue when it registers a subscription.

Data type	String
Units	En_US ASCII characters
Default	Null
Range	1 through 48 ASCII characters

Broker CC subscription queue:

The name of the broker's queue from which non-durable subscription messages are retrieved for a ConnectionConsumer. This property applies only for use of the Web container.

The name of the broker's queue from which non-durable subscription messages are retrieved for a ConnectionConsumer. This property applies only for use of the Web container.

Data type	String
Units	En_US ASCII characters
Default	Null
Range	1 through 48 ASCII characters

Broker version:

Whether the message broker is provided by the WebSphere MQ MA0C Supportpac or newer versions of WebSphere message broker products

Data type	Enum
Units	Not applicable
Default	Advanced
Range	Advanced The message broker is provided by newer versions of WebSphere message broker products, such as WebsSphere MQ Integrator and EventBroker. Basic The message broker is provided by the WebSphere MQ MA0C SupportPac (MQSeries - Publish/Subscribe) or MQSI working in MA0C compatibility mode.

Model queue definition:

The name of the model queue definition that the broker can use to create dynamic queues for non-default streams if the stream queue does not already exist

The name of the model queue definition that the broker can use to create dynamic queues to receive publications for streams other than the default stream. This is only used if the stream queue does not already exist. If this model queue definition does not exist, all stream queues must be defined by the administrator.

Data type	String
Units	En_US ASCII characters
Default	Null
Range	1 through 48 ASCII characters

CCSID:

The coded character set identifier for use with the WebSphere MQ queue manager.

This coded character set identifier (CCSID) must be one of the CCSIDs supported by WebSphere MQ.

Data type	String
Units	Integer
Default	Null
Range	1 through 65535

For more information about supported CCSIDs, and about converting between message data from one coded character set to another, see the *WebSphere MQ System Administration* and the *WebSphere MQ Application Programming Reference* books. These are available from the WebSphere MQ messaging multiplatform and platform-specific books Web pages; for example, at <http://www-3.ibm.com/software/ts/mqseries/library/manualsa/manuals/platspecific.html>, the IBM Publications Center, or from the WebSphere MQ collection kit, SK2T-0730.

Clone Support:

Select this checkbox to enable clone support to allow the same durable subscription across topic clones.

Data type	Enum
Units	Not applicable

Default Range

Cleared
Selected Clone support is enabled.
Cleared Clone support is disabled.

If you select this property, you must also specify a value for the **Client ID** property.

Client ID:

The JMS client identifier used for connections to the queue manager.

Data type Range

String
A valid JMS client ID, as ASCII characters

XA Enabled:

Specifies whether the connection factory is for XA or non-XA coordination of messages and controls if the application server uses XA QCF/TCF. Enable XA if multiple resources are not used in the same transaction.

If you set this property to `NON_XA`, the JMS session is still enlisted in a transaction, but uses the resource manager local transaction calls (`session.commit` and `session.rollback`) instead of XA calls. This can lead to an improvement in performance. However, this means that only a single resource can be enlisted in a transaction in WebSphere Application Server.

5.1+ In WBI Server Foundation the last participant support enables you to enlist one non-XA resource with other XA-capable resources.

For a WebSphere Topic Connection Factory with the **Port** property set to `DIRECT` this property does not apply, and always takes the value `NON_XA`.

Data type Units Default Range

Enum
Not applicable
XA enabled
XA The connection factory is for XA-coordination of messages
NON_XA The connection factory is for non-XA coordination of messages
Do not enable XA when the message queue or topic received is the only resource in the transaction. Enable XA when other resources, including other queues or topics, are involved.

Recommended

Connection pool:

Specifies an optional set of connection pool settings.

Connection pool properties are common to all J2C connectors.

The application server pools connections and sessions with the JMS provider to improve performance. This is independent from any WebSphere MQ connection pooling. You need to configure the connection and session pool properties appropriately for your applications, otherwise you may not get the connection and session behavior that you want.

Change the size of the connection pool if concurrent server-side access to the JMS resource exceeds the default value. The size of the connection pool is set on a per queue or topic basis. See the following table for details.

Session pool:

An optional set of session pool settings.

This link provides a panel of optional connection pool properties, common to all J2C connectors.

The application server pools connections and sessions with the JMS provider to improve performance. This is independent from any WebSphere MQ connection pooling. You need to configure the connection and session pool properties appropriately for your applications, otherwise you may not get the connection and session behavior that you want.

Custom properties:

An optional set of name and value pairs for custom properties passed to WebSphere MQ.

You can use the Custom properties page to define the following WebSphere MQ properties. These properties apply only if you set the property **Transport type=CLIENT**.

- SSLCIPHERSUITE
- SSLCRL
- SSLPEERNAME
- MSGSELECTION
- SUBSTORE

WebSphere MQ topic connection factory custom properties:

Use this panel to view or change an optional set of name and value pairs for custom properties of the selected topic connection factory for use with the WebSphere MQ JMS provider. These custom properties are passed to WebSphere MQ.

To view this administrative console page, click **Resources-> WebSphere MQ JMS Provider-> (In content pane, under Additional Properties) WebSphere MQ Topic Connection Factories-> connection_factory-> Custom Properties**

You can use the Custom properties page to define the following WebSphere MQ properties. These properties apply only if you set the property **Transport type=CLIENT**.

- SSLCIPHERSUITE
- SSLCRL
- SSLPEERNAME
- MSGSELECTION
- SUBSTORE

For more information about setting the SSL properties for WebSphere MQ, see the section SSL properties in the *WebSphere MQ Using Java* book.

SSLCIPHERSUITE:

The cipher suite to use for SSL connection.

Set this property to a valid cipher suite provided by your JSSE provider; it must match the CipherSpec named on the SVRCONN channel named by the **Channel** property.

You must set this property if the SSLPEERNAME is to be set.

SSLCRL:

A list of zero or more CRL (Certificate Revocation List) servers used to check for SSL certificate revocation. (Use of this property requires a WebSphere MQ JVM at Java 2 version 1.4.)

The value is a space-delimited list of entries of the form:

```
ldap://hostname:[port]
```

optionally followed by a single / (forward slash). If *port* is omitted, the default LDAP port of 389 is assumed. At connect-time, the SSL certificate presented by the server is checked against the specified CRL servers. For more information about CRL security, see the section “Working with Certificate Revocation Lists” in the *WebSphere MQ Security book*; for example at: <http://publibfp.boulder.ibm.com/epubs/html/csqzas01/csqzas012w.htm#IDX2254>.

SSLPEERNAME:

For SSL, a *distinguished name* skeleton that must match the name provided by the WebSphere MQ queue manager. The distinguished name is used to check the identifying certificate presented by the server at connect-time.

If SSLPEERNAME is not set, such checking is performed. SSLPEERNAME is ignored if SSLCIPHERSUITE is not specified.

The SSLPEERNAME property is a list of attribute name and value pairs separated by commas or semicolons. For example:

```
SSLPEERNAME(CN=QMGR.*, OU=IBM, OU=WEBSHERE)
```

The example given checks the identifying certificate presented by the server at connect-time. For the connection to succeed, the certificate must have a Common Name beginning QMGR., and must have at least two Organizational Unit names, the first of which is IBM and the second WEBSHERE. Checking is not case-sensitive.

For more details about distinguished names and their use with WebSphere MQ, see the *WebSphere MQ Security book*; for example, the section “Distinguished Names” at <http://publibfp.boulder.ibm.com/epubs/html/csqzas01/csqzas010p.htm#HDRDCDN>.

MSGSELECTION:

Determines whether message selection is done by the JMS Client or by the Broker.

If you set the property **Transport type=BINDINGS**, message selection is always done by the Broker and the setting of this property is ignored.

Data type	String
Default	CLIENT

Range

CLIENT

Message selection is done by the JMS Client.

BROKER

Message selection is done by the Broker.

SUBSTORE:

Where WebSphere MQ JMS should store persistent data relating to active subscriptions.

Data type

String

Default

MIGRATE

Range

BROKER

Subscription information is stored by the publish/subscribe broker used by the application.

If a non-durable subscriber fails, the subscription is deregistered from the broker as soon as possible. The broker adds a response to this deregistration onto the SYSTEM.JMS.REPORT.QUEUE, which is used to clean up after the failed subscriber. With SUBSTORE(BROKER), a separate cleanup thread is run regularly in the background of each JMS publish/subscribe application.

MIGRATE

This option dynamically selects the queue-based or broker-based subscription store based on the levels of queue manager and publish/subscribe broker installed. If both queue manager and broker are capable of supporting SUBSTORE(BROKER), this behaves as SUBSTORE(BROKER); otherwise it behaves as SUBSTORE(QUEUE). Additionally, SUBSTORE(MIGRATE) transfers durable subscription information from the queue-based subscription store to the broker-based store.

This provides an easy migration path from older versions of WebSphere MQ JMS, WebSphere MQ, and publish/subscribe broker.

QUEUE

Subscription information is stored on SYSTEM.JMS.ADMIN.QUEUE and SYSTEM.JMS.PS.STATUS.QUEUE on the local queue manager. SUBSTORE(QUEUE) is provided for compatibility with versions of MQSeries JMS.

For more information about the use of this property in WebSphere MQ, see the section Subscription stores in the *WebSphere MQ Using Java* book.

WebSphere MQ queue destination collection:

The queue destinations configured in the WebSphere MQ JMS provider for point-to-point messaging with JMS queues.

This panel shows a list of the WebSphere MQ queue destinations with a summary of their configuration properties.

To view this administrative console page, use the administrative console to complete the following steps:

1. In the navigation pane, expand **Resources-> WebSphere MQ JMS provider**.
2. In the content pane, select the JMS provider that you want to support the queue destination.
3. In the Additional Properties list in the contents pane, select **WebSphere MQ Queue Destination**.

To view or change the properties of a queue destination, select its name in the list displayed.

To act on one or more of the queue destinations listed, click the check boxes next to the names of the objects that you want to act on, then use the buttons provided.

WebSphere MQ queue settings:

Use this panel to view or change the configuration properties of the selected queue destination for use with the WebSphere MQ JMS provider.

A queue destination is used to configure the properties of a JMS queue. Connections to the queue are created by the associated queue connection factory for the WebSphere MQ JMS provider. A queue for use with the WebSphere MQ JMS provider has the following properties.

Note:

- The property values that you specify must match the values that you specified when configuring WebSphere MQ JMS resources. For more information about configuring WebSphere MQ JMS resources, see the *WebSphere MQ Using Java* book.
- In WebSphere MQ, names can have a maximum of 48 characters, with the exception of channels which have a maximum of 20 characters.

To view this administrative console page, click **Resources-> WebSphere MQ JMS Provider-> (In content pane, under Additional Properties) WebSphere MQ Queue Destinations-> *destination_name***

Name:

The name by which the queue is known for administrative purposes within IBM WebSphere Application Server.

Data type String

JNDI name:

The JNDI name that is used to bind the queue into the application server's name space.

As a convention, use the fully qualified JNDI name; for example, in the form *jms/Name*, where *Name* is the logical name of the resource.

This name is used to link the platform binding information. The binding associates the resources defined by the deployment descriptor of the module to the actual (physical) resources bound into JNDI by the platform.

Data type String

Description:

A description of the queue, for administrative purposes

Data type String
Default Null

Category:

A category used to classify or group this queue, for your IBM WebSphere Application Server administrative records.

Data type String

Persistence:

Whether all messages sent to the destination are persistent, non-persistent, or have their persistence defined by the application

Data type Enum
Units Not applicable
Default APPLICATION_DEFINED
Range

- Application defined**
Messages on the destination have their persistence defined by the application that put them onto the queue.
- Queue defined**
[WebSphere MQ destination only] Messages on the destination have their persistence defined by the WebSphere MQ queue definition properties.
- Persistent**
Messages on the destination are persistent.
- Non persistent**
Messages on the destination are not persistent.

Priority:

Whether the message priority for this destination is defined by the application or the **Specified priority** property

Data type Enum
Units Not applicable
Default APPLICATION_DEFINED
Range

- Application defined**
The priority of messages on this destination is defined by the application that put them onto the destination.
- Queue defined**
[WebSphere MQ destination only] Messages on the destination have their persistence defined by the WebSphere MQ queue definition properties.
- Specified**
The priority of messages on this destination is defined by the **Specified priority** property. *If you select this option, you must define a priority on the **Specified priority** property.*

Specified priority:

If the **Priority** property is set to *Specified*, type here the message priority for this queue, in the range 0 (lowest) through 9 (highest)

If the **Priority** property is set to *Specified*, messages sent to this queue have the priority value specified by this property.

Data type	Integer
Units	Message priority level
Default	Null
Range	0 (lowest priority) through 9 (highest priority)

Expiry:

Whether the expiry timeout for this queue is defined by the application or the **Specified expiry** property, or messages on the queue never expire (have an unlimited expiry timeout)

Data type	Enum
Units	Not applicable
Default	APPLICATION_DEFINED
Range	Application defined The expiry timeout for messages on this queue is defined by the application that put them onto the queue. Specified The expiry timeout for messages on this queue is defined by the Specified expiry property. <i>If you select this option, you must define a timeout on the Specified expiry property.</i> Unlimited Messages on this queue have no expiry timeout, so those messages never expire.

Specified expiry:

If the **Expiry timeout** property is set to *Specified*, type here the number of milliseconds (greater than 0) after which messages on this queue expire

Data type	Integer
Units	Milliseconds
Default	Null
Range	Greater than or equal to 0 • 0 indicates that messages never timeout • Other values are an integer number of milliseconds

Base queue name:

The name of the queue to which messages are sent, on the queue manager specified by the **Base queue manager name** property

Data type	String
Default	Null
Range	1 through 48 ASCII characters

Base queue manager name:

The name of the WebSphere MQ queue manager to which messages are sent

This queue manager provides the queue specified by the **Base queue name** property.

Data type	String
Units	En_US ASCII characters
Default	Null
Range	A valid WebSphere MQ Queue Manager name, as 1 through 48 ASCII characters

CCSID:

The coded character set identifier for use with the WebSphere MQ queue manager.

This coded character set identifier (CCSID) must be one of the CCSIDs supported by WebSphere MQ.

Data type	String
Units	Integer
Default	Null
Range	1 through 65535

For more information about supported CCSIDs, and about converting between message data from one coded character set to another, see the *WebSphere MQ System Administration* and the *WebSphere MQ Application Programming Reference* books. These are available from the WebSphere MQ messaging multiplatform and platform-specific books Web pages; for example, at <http://www-3.ibm.com/software/ts/mqseries/library/manualsa/manuals/platspecific.html>, the IBM Publications Center, or from the WebSphere MQ collection kit, SK2T-0730.

Use native encoding:

Select this checkbox to indicate that the queue destination should use native encoding (appropriate encoding values for the Java platform).

Data type	Enum
Units	Not applicable
Default	Cleared
Range	Cleared Native encoding is not used, so specify the properties below for integer, decimal, and floating point encoding. Selected Native encoding is used (to provide appropriate encoding values for the Java platform). For more information about encoding properties, see the WebSphere MQ <i>Using Java</i> document.

Integer encoding:

If native encoding is not enabled, select whether integer encoding is normal or reversed.

Data type	Enum
Units	Not applicable
Default	NORMAL

Range**NORMAL**

Normal integer encoding is used.

REVERSED

Reversed integer encoding is used.

For more information about encoding properties, see the WebSphere MQ *Using Java* document.

Decimal encoding:

If native encoding is not enabled, select whether decimal encoding is normal or reversed.

Data type

Enum

Units

Not applicable

Default

NORMAL

Range**NORMAL**

Normal decimal encoding is used.

REVERSED

Reversed decimal encoding is used.

For more information about encoding properties, see the WebSphere MQ *Using Java* document.

Floating point encoding:

If native encoding is not enabled, select the type of floating point encoding.

Data type

Enum

Units

Not applicable

Default

IEEEENORMAL

Range**IEEEENORMAL**

IEEE normal floating point encoding is used.

IEEEEVERSED

IEEE reversed floating point encoding is used.

S390 S390 floating point encoding is used.

For more information about encoding properties, see the WebSphere MQ *Using Java* document.

Target client:

Whether the receiving application is JMS-compliant or is a traditional WebSphere MQ application

Data type

Enum

Units

Not applicable

Default

MQSeries

Range**MQSeries**

The target is a non-JMS, traditional WebSphere MQ application.

JMS The target is a JMS-compliant application.*Queue manager host:*

The name of host for the queue manager on which the queue destination is created.

Data type

String

Default	Null
Range	A valid TCP/IP hostname

Queue manager port:

The number of the port used by the queue manager on which this queue is defined.

Data type	String
Units	A valid TCP/IP port number.
Default	Null
Range	A valid TCP/IP port number. This port must be configured on the WebSphere MQ queue manager.

Server connection channel name:

The name of the channel used for connection to the WebSphere MQ queue manager.

Data type	String
Default	Null
Range	1 through 20 ASCII characters

User name:

The user ID used, with the **Password** property, for authentication when connecting to the queue manager to define the queue destination.

If you specify a value for the **User name** property, you must also specify a value for the **Password** property.

Data type	String
Default	Null

Password:

The password, used with the **User name** property, for authentication when connecting to the queue manager to define the queue destination.

If you specify a value for the **User name** property, you must also specify a value for the **Password** property.

Data type	String
Default	Null

WebSphere MQ queue settings (MQ Config):

Use this panel to view or change the configuration properties defined to WebSphere MQ for the selected queue destination.

A queue destination is used to configure the properties of a JMS queue. A queue for use with the WebSphere MQ JMS provider has the following extra properties defined to WebSphere MQ.

Notes

Note:

- Some properties displayed are read-only and cannot be changed.
- These MQ Config properties can be used only to view or change the properties of local queues. You cannot use MQ Config to administer alias or remote queues.
- To be able to view or change properties, the WebSphere MQ Queue Manager on which the queue resides must be configured for remote administration and be running.
- The property values that you specify must match the values that you specified when configuring WebSphere MQ JMS resources. For more information about configuring WebSphere MQ JMS resources, see the *WebSphere MQ: Using Java* book; for example from the WebSphere MQ multiplatform library Web page at <http://www.ibm.com/software/ts/mqseries/library/manualsa/manuals/crosslatest.html>.
- In WebSphere MQ, names can have a maximum of 48 characters, with the exception of channels which have a maximum of 20 characters.

To view this administrative console page, click **Resources-> WebSphere MQ JMS Provider-> (In content pane, under Additional Properties) WebSphere MQ Queue Destinations-> destination_name-> (In content pane, under Additional Properties) MQ Config**

Base queue name:

The name of the local queue to which messages are sent, on the queue manager specified by the **Base queue manager name** property.

Data type String

Base queue manager name:

The name of the WebSphere MQ queue manager to which messages are sent.

This queue manager provides the queue specified by the **Base queue name** property.

Data type String

Queue manager host:

The name of host for the queue manager on which the queue destination is created.

Data type String

Queue manager port:

The number of the port used by the queue manager on which this queue is defined.

Data type Integer
Range A valid TCP/IP port number. This port must be configured on the WebSphere MQ queue manager.

Server connection channel name:

The name of the channel used for connection to the WebSphere MQ queue manager.

Data type String
Range 1 through 20 ASCII characters

User ID:

The user ID used, with the **Password** property, for authentication when connecting to the queue manager to define the queue destination.

If you specify a value for the **User name** property, you must also specify a value for the **Password** property.

Data type String

Password:

The password, used with the **User name** property, for authentication when connecting to the queue manager to define the queue destination.

If you specify a value for the **User name** property, you must also specify a value for the **Password** property.

Data type String

Name:

The name of the queue defined to the WebSphere MQ queue manager.

Data type String
Range 1 through 48 ASCII characters.

Description:

The WebSphere MQ queue description, for administrative purposes within WebSphere MQ.

Data type String
Default Null
Range 1 through 64 ASCII characters.

Inhibit Put:

Whether or not put operations are allowed for this queue.

Data type Enum
Units Not applicable
Default Cleared
Range **Allowed** Put operations are allowed for this queue.
Not allowed Put operations are not allowed for this queue.

Persistence:

Whether all messages sent to the destination are persistent, non-persistent, or have their persistence defined by the application

Data type Enum
Units Not applicable

Default
Range

APPLICATION_DEFINED

Application defined

Messages on the destination have their persistence defined by the application that put them onto the queue.

Queue defined

[WebSphere MQ destination only] Messages on the destination have their persistence defined by the WebSphere MQ queue definition properties.

Persistent

Messages on the destination are persistent.

Non persistent

Messages on the destination are not persistent.

Cluster name:

The name of the cluster to which the WebSphere MQ queue manager belongs.

If you specify a value for **Cluster name**, you cannot specify a value for **Cluster name list**. Cluster names must conform to the rules described in the *WebSphere MQ MQSC Command Reference* book.

Data type

String

Default

Null

Range

A valid WebSphere MQ name for a queue manager cluster, as 1 through 48 ASCII characters

Cluster name list:

The name of the cluster namelist to which the WebSphere MQ queue manager belongs.

If you specify a value for **Cluster name**, you cannot specify a value for **Cluster name list**.

Data type

String

Default

Null

Range

A valid WebSphere MQ name for a list of queue manager clusters, as 1 through 48 ASCII characters

Default Binding:

The default binding to be used when the queue is defined as a cluster queue.

Data type

Enum

Default

Null

Range

A

On open

The queue handle is bound to a specific instance of the cluster queue when the queue is opened.

Not fixed

The queue handle is not bound to any particular instance of the cluster queue. This allows the queue manager to select a specific queue instance when the message is put, and to change that selection subsequently should the need arise.

Inhibit Get:

Whether or not get operations are allowed for this queue.

Data type	Enum
Units	Not applicable
Default	Cleared
Range	Allowed Get operations are allowed for this queue. Not allowed Get operations are not allowed for this queue.

Maximum queue depth:

The maximum number of messages allowed on the queue.

Data type	Integer
Units	Messages
Default	
Range	A value greater than or equal to zero, and less than or equal to: <ul style="list-style-type: none">• 999 999 999 if the queue is on OS/390• 640 000 if the queue is on any other WebSphere MQ platform For more information about the maximum value allowed, see the <i>WebSphere MQ MQSC Command Reference</i> . If this value is reduced, any message that is already on the queue are not affected, even if the number of messages exceeds the new maximum.

Maximum Message Length:

The maximum length, in bytes, of messages on this queue.

Data type	Integer
Units	Bytes
Default	
Range	A value greater than or equal to zero, and less than or equal to the maximum message length for the queue manager and WebSphere MQ platform. For more information about the maximum value allowed, see the <i>WebSphere MQ MQSC Command Reference</i> . If this value is reduced, any message that is already on the queue are not affected, even if the message length exceeds the new maximum.

Shareability:

Whether multiple applications can get messages from this queue.

Data type	Enum
Units	Not applicable
Default	Not shareable

Range**Not shareable**

Only one application instance can get messages from the queue.

Shareable

More than one application instance can get messages from the queue.

Input Open Option:

The default share option for applications opening this queue for input

Data type

Enum

Units

Not applicable

Default

Cleared

Range**Exclusive**

The open request is for exclusive input from the queue.

Shared

The open request is for shared input from the queue.

Message Delivery Sequence:

The order in which messages are delivered from the queue in response to get requests.

Data type

Enum

Units

Not applicable

Default

Priority

Range**Priority**

Messages are delivered in first-in-first-out (FIFO) order within priority. This is the default supplied with WebSphere MQ, but your installation might have changed it.

FIFO

Messages are delivered in FIFO order. Priority is ignored for messages on this queue.

Backout threshold:

The maximum number of times that a message can be backed out. If this threshold is reached, the message is requeued on the backout queue specified by the **Backout Requeue name** property.

The WebSphere MQ queue manager keeps a record of the number of times that each message has been backed out. When this number reaches a configurable threshold, the connection consumer requeues the message on a named backout queue. If this requeue fails for any reason, the message is removed from the queue and either requeued to the dead-letter queue, or discarded.

Data type

Integer

Default

0

Range

0 Never requeue messages

1 or more

The number of times that a message has been backed, at which the message is requeued on a named backout queue.

Backout Requeue name:

The name of the backout queue to which messages are requeued if they have been backed out more than the backout threshold.

The WebSphere MQ queue manager keeps a record of the number of times that each message has been backed out. When this number reaches a configurable threshold, the connection consumer requeues the message on a named backout queue. If this requeue fails for any reason, the message is removed from the queue and either requeued to the dead-letter queue, or discarded.

Data type	String
Default	Null
Range	1 through 48 characters.

Harden Get Backout:

Whether hardening should be used to ensure that the count of the number of times that a message has been backed out is accurate.

Data type	Enum
Units	Not applicable
Default	Cleared
Range	Not hardened The count is not hardened. This is the default supplied with WebSphere MQ, but your installation might have changed it. Hardened The count is hardened.

WebSphere MQ topic destination collection:

The topic destinations configured in the WebSphere MQ JMS provider for publish/subscribe messaging with JMS topics. Use this panel to create or delete topic destinations, or to select a topic destination to view or change its configuration properties.

This panel shows a list of the WebSphere MQ topic destinations with a summary of their configuration properties.

To view this administrative console page, use the administrative console to complete the following steps:

1. In the navigation pane, expand **Resources-> WebSphere MQ JMS provider**.
2. In the content pane, select the JMS provider that you want to support the topic destination.
3. In the Additional Properties list in the contents pane, select **WebSphere MQ Topic Destination**.

To view or change the properties of a topic destination, select its name in the list displayed.

To act on one or more of the topic destinations listed, click the check boxes next to the names of the objects that you want to act on, then use the buttons provided.

WebSphere MQ topic settings:

Use this panel to view or change the configuration properties of the selected topic destination for use with the WebSphere MQ JMS provider.

A topic destination is used to configure the properties of a JMS topic for the associated JMS provider. Connections to the topic are created by the associated topic connection factory. A topic for use with the WebSphere MQ JMS provider has the following properties.

Note:

- The property values that you specify must match the values that you specified when configuring WebSphere MQ JMS resources. For more information about configuring WebSphere MQ JMS resources, see the WebSphere MQ *Using Java* book.
- In WebSphere MQ, names can have a maximum of 48 characters, with the exception of channels which have a maximum of 20 characters.

To view this administrative console page, click **Resources-> WebSphere MQ JMS Provider-> (In content pane, under Additional Properties) WebSphere MQ Topic Destinations-> destination_name**

Name:

The name by which the topic is known for administrative purposes within IBM WebSphere Application Server.

Data type String

JNDI name:

The JNDI name that is used to bind the topic into the application server's name space.

As a convention, use the fully qualified JNDI name; for example, in the form `jms/Name`, where *Name* is the logical name of the resource.

This name is used to link the platform binding information. The binding associates the resources defined by the deployment descriptor of the module to the actual (physical) resources bound into JNDI by the platform.

Data type String

Description:

A description of the topic, for administrative purposes within IBM WebSphere Application Server.

Data type String
Default Null

Category:

A category used to classify or group this topic, for your IBM WebSphere Application Server administrative records.

Data type String

Persistence:

Whether all messages sent to the destination are persistent, non-persistent, or have their persistence defined by the application

Data type Enum
Units Not applicable
Default APPLICATION_DEFINED

Range

Application defined

Messages on the destination have their persistence defined by the application that put them onto the queue.

Queue defined

[WebSphere MQ destination only] Messages on the destination have their persistence defined by the WebSphere MQ queue definition properties.

Persistent

Messages on the destination are persistent.

Non persistent

Messages on the destination are not persistent.

Priority:

Whether the message priority for this destination is defined by the application or the **Specified priority** property

Data type
Units
Default
Range

Enum
Not applicable
APPLICATION_DEFINED

Application defined

The priority of messages on this destination is defined by the application that put them onto the destination.

Queue defined

[WebSphere MQ destination only] Messages on the destination have their persistence defined by the WebSphere MQ queue definition properties.

Specified

The priority of messages on this destination is defined by the **Specified priority** property. *If you select this option, you must define a priority on the **Specified priority** property.*

Specified priority:

If the **Priority** property is set to Specified, type here the message priority for this queue, in the range 0 (lowest) through 9 (highest)

If the **Priority** property is set to Specified, messages sent to this queue have the priority value specified by this property.

Data type
Units
Default
Range

Integer
Message priority level
Null
0 (lowest priority) through 9 (highest priority)

Expiry:

Whether the expiry timeout for this queue is defined by the application or the **Specified expiry** property, or messages on the queue never expire (have an unlimited expiry timeout)

Data type
Units
Default

Enum
Not applicable
APPLICATION_DEFINED

Range**Application defined**

The expiry timeout for messages on this queue is defined by the application that put them onto the queue.

Specified

The expiry timeout for messages on this queue is defined by the **Specified expiry** property. *If you select this option, you must define a timeout on the **Specified expiry** property.*

Unlimited

Messages on this queue have no expiry timeout, so those messages never expire.

Specified expiry:

If the **Expiry timeout** property is set to Specified, type here the number of milliseconds (greater than 0) after which messages on this queue expire

Data type

Integer

Units

Milliseconds

Default

Null

Range

Greater than or equal to 0

- 0 indicates that messages never timeout
- Other values are an integer number of milliseconds

Base topic name:

The name of the topic to which messages are sent

Data type

String

Range

Depends on the broker used. For details, see the documentation for your broker; for example the *WebSphere MQ Event Broker* library at <http://www-3.ibm.com/software/ts/mqseries/library/manualsa/manuals/wsmqebv21.html>.

CCSID:

The coded character set identifier for use with the WebSphere MQ queue manager.

This coded character set identifier (CCSID) must be one of the CCSIDs supported by WebSphere MQ.

Data type

String

Units

Integer

Default

Null

Range

1 through 65535

For more information about supported CCSIDs, and about converting between message data from one coded character set to another, see the *WebSphere MQ System Administration* and the *WebSphere MQ Application Programming Reference* books. These are available from the WebSphere MQ messaging multiplatform and platform-specific books Web pages; for example, at <http://www-3.ibm.com/software/ts/mqseries/library/manualsa/manuals/platspecific.html>, the IBM Publications Center, or from the WebSphere MQ collection kit, SK2T-0730.

Use native encoding:

Select this checkbox to indicate that the queue destination should use native encoding (appropriate encoding values for the Java platform).

Data type	Enum
Units	Not applicable
Default	Cleared
Range	Cleared Native encoding is not used, so specify the properties below for integer, decimal, and floating point encoding. Selected Native encoding is used (to provide appropriate encoding values for the Java platform).

For more information about encoding properties, see the WebSphere MQ *Using Java* document.

Integer encoding:

If native encoding is not enabled, select whether integer encoding is normal or reversed.

Data type	Enum
Units	Not applicable
Default	NORMAL
Range	NORMAL Normal integer encoding is used. REVERSED Reversed integer encoding is used.

For more information about encoding properties, see the WebSphere MQ *Using Java* document.

Decimal encoding:

If native encoding is not enabled, select whether decimal encoding is normal or reversed.

Data type	Enum
Units	Not applicable
Default	NORMAL
Range	NORMAL Normal decimal encoding is used. REVERSED Reversed decimal encoding is used.

For more information about encoding properties, see the WebSphere MQ *Using Java* document.

Floating point encoding:

If native encoding is not enabled, select the type of floating point encoding.

Data type	Enum
Units	Not applicable
Default	IEEE NORMAL

Range**IEEEENORMAL**

IEEE normal floating point encoding is used.

IEEEEVERSED

IEEE reversed floating point encoding is used.

S390 S390 floating point encoding is used.For more information about encoding properties, see the *WebSphere MQ Using Java* document.*Target client type:*

Whether the receiving application is JMS-compliant or is a traditional WebSphere MQ application

Data type

Enum

Units

Not applicable

Default

MQ

Range**MQ** The target is a non-JMS, traditional WebSphere MQ application.**JMS** The target is a JMS-compliant application.*Broker Dur Sub Queue:*

The name of the broker's queue from which durable subscription messages are retrieved

The name of the broker's queue from which durable subscription messages are retrieved. The subscriber specifies the name of the queue when it registers a subscription.

Data type

String

Units

En_US ASCII characters

Default

Null

Range

1 through 48 ASCII characters

Broker CC Dur Sub Queue:

The name of the broker's queue from which durable subscription messages are retrieved for a ConnectionConsumer. This property applies only for use of the Web container.

The name of the broker's queue from which durable subscription messages are retrieved for a ConnectionConsumer. This property applies only for use of the Web container.

Data type

String

Units

En_US ASCII characters

Default

Null

Range

1 through 48 ASCII characters

JMS connection factory collection:

The JMS connection factories configured in the associated JMS provider for both point-to-point and publish/subscribe messaging. Use this panel to create or delete JMS connection factories, or to select a connection factory to view or change its configuration properties.

This panel shows a list of the generic JMS connection factories with a summary of their configuration properties.

To view this administrative console page, use the administrative console to complete the following steps:

1. In the navigation pane, expand **Resources-> Generic JMS Providers**.
2. In the content pane, select the JMS provider that you want to support the JMS connection factory.
3. In the Additional Properties list in the contents pane, select **Generic JMS Connection Factory**.

To view or change the properties of a JMS connection factory, select its name in the list displayed.

To act on one or more of the JMS connection factories listed, click the check boxes next to the names of the objects that you want to act on, then use the buttons provided.

Generic JMS connection factory settings:

Use this panel to view or change the configuration properties of the selected JMS connection factory for use with the associated JMS provider. These configuration properties control how connections are created to the associated JMS destination.

A JMS connection factory is used to create connections to JMS destinations. The JMS connection factory is created by the associated JMS provider. A JMS connection factory for a generic JMS provider (other than the embedded WebSphere JMS provider or the WebSphere MQ JMS provider) has the following properties:

To view this administrative console page, click **Resources-> Generic JMS Providers-> *provider_name*-> JMS Connection Factories-> *connection_factory***

Scope:

Specifies the level to which this resource definition is visible -- the cell, node, or server level.

Resources such as JMS Providers, Namespace bindings, or shared libraries can be defined at multiple scopes, with resources defined at more specific scopes overriding duplicates which are defined at more general scopes.

Note that no matter what the scope of a defined resource, the resource's properties only apply at an individual server level. For example, if you define the scope of a data source at the Cell level, all users in that Cell can look up and use that data source, which is unique within that Cell. However, resource property settings are local to each server in the Cell. For example, if you set **Max Connections** to 10, then each server in that Cell can have 10 connections.

Cell The most general scope. Resources defined at the Cell scope are visible from all Nodes and servers, unless they are overridden. To view resources defined in the cell scope, do not specify a server or a node name in the scope selection form.

Node The default scope for most resource types. Resources defined at the Node scope override any duplicates defined at the Cell scope and are visible to all servers on the same node, unless they are overridden at a server scope on that node. To view resources defined in a node scope, do not specify a server, but select a node name in the scope selection form.

Server The most specific scope for defining resources. Resources defined at the Server scope override any duplicate resource definitions defined at the Cell scope or parent Node scope and are visible only to a specific server. To view resources defined in a server scope, specify a server name as well as a node name in the scope selection form.

When resources are created, they are always created into the current scope selected in the panel. To view resources in other scopes, specify a different node or server in the scope selection form.

Data type String

Name:

The name by which this JMS connection factory is known for administrative purposes within IBM WebSphere Application Server. The name must be unique within the associated JMS provider.

Data type String

Type:

Whether this connection factory is for creating JMS queue destinations or JMS topic destinations.

Select one of the following options:

Queue

A JMS queue connection factory for point-to-point messaging.

Topic A JMS topic connection factory for publish/subscribe messaging.

JNDI name:

The JNDI name that is used to bind the connection factory into the application server's name space.

As a convention, use the fully qualified JNDI name; for example, in the form *jms/Name*, where *Name* is the logical name of the resource.

This name is used to link the platform binding information. The binding associates the resources defined by the deployment descriptor of the module to the actual (physical) resources bound into JNDI by the platform.

Data type String

Description:

A description of this connection factory for administrative purposes within IBM WebSphere Application Server.

Data type String

Default Null

Category:

A category used to classify or group this connection factory, for your IBM WebSphere Application Server administrative records.

Data type String

Component-managed Authentication Alias:

This alias specifies a user ID and password to be used to authenticate connection to a JMS provider for application-managed authentication.

This property provides a list of the J2C authentication data entry aliases that have been defined to WebSphere Application Server. You can select a data entry alias to be used to authenticate the creation of a new connection to the JMS provider.

If you have enabled global security for WebSphere Application Server, select the alias that specifies the user ID and password used to authenticate the creation of a new connection to the JMS provider. The use

of this alias depends on the resource authentication (res-auth) setting declared in the connection factory resource reference of an application component's deployment descriptors.

Note: User IDs longer than 12 characters cannot be used for authentication with the embedded WebSphere JMS provider. For example, the default Windows NT user ID, **Administrator**, is not valid for use with embedded WebSphere messaging, because it contains 13 characters. Therefore, an authentication alias for a WebSphere JMS provider connection factory must specify a user ID no longer than 12 characters.

Container-managed Authentication Alias:

This alias specifies a user ID and password to be used to authenticate connection to a JMS provider for container-managed authentication.

This property provides a list of the J2C authentication data entry aliases that have been defined to WebSphere Application Server. You can select a data entry alias to be used to authenticate the creation of a new connection to the JMS provider.

If you have enabled global security for WebSphere Application Server, select the alias that specifies the user ID and password used to authenticate the creation of a new connection to the JMS provider. The use of this alias depends on the resource authentication (res-auth) setting declared in the connection factory resource reference of an application component's deployment descriptors.

Note: User IDs longer than 12 characters cannot be used for authentication with the embedded WebSphere JMS provider. For example, the default Windows NT user ID, **Administrator**, is not valid for use with embedded WebSphere messaging, because it contains 13 characters. Therefore, an authentication alias for a WebSphere JMS provider connection factory must specify a user ID no longer than 12 characters.

Mapping-Configuration Alias:

Allows users to select from the **Security > JAAS Configuration > Application Logins Configuration** list.

The **DefaultPrincipalMapping** JAAS configuration maps the authentication alias to the userid and password. You may define and use other mapping configurations. For more information about the mapping configurations, see "Java Authentication and Authorization service configuration entry settings" in the Information center.

Data type Pick-list

Connection pool:

Specifies an optional set of connection pool settings.

Connection pool properties are common to all J2C connectors.

The application server pools connections and sessions with the JMS provider to improve performance. This is independent from any WebSphere MQ connection pooling. You need to configure the connection and session pool properties appropriately for your applications, otherwise you may not get the connection and session behavior that you want.

Change the size of the connection pool if concurrent server-side access to the JMS resource exceeds the default value. The size of the connection pool is set on a per queue or topic basis. See the following table for details.

Session pool:

An optional set of session pool settings.

This link provides a panel of optional connection pool properties, common to all J2C connectors.

The application server pools connections and sessions with the JMS provider to improve performance. This is independent from any WebSphere MQ connection pooling. You need to configure the connection and session pool properties appropriately for your applications, otherwise you may not get the connection and session behavior that you want.

Custom properties:

An optional set of name and value pairs for custom properties passed to the JMS provider.

Generic JMS destination collection:

The JMS destinations configured in the associated JMS provider for point-to-point and publish/subscribe messaging. Use this panel to create or delete JMS destinations, or to select a JMS destination to view or change its configuration properties.

To view this administrative console page, use the administrative console to complete the following steps:

1. In the navigation pane, expand **Resources-> Generic JMS Providers**.
2. In the content pane, select the JMS provider that you want to support the JMS destination.
3. In the Additional Properties list in the contents pane, select **Generic JMS Destination**.

To view or change the properties of a JMS destination, select its name in the list displayed.

To act on one or more of the JMS destinations listed, click the check boxes next to the names of the objects that you want to act on, then use the buttons provided.

Generic JMS destination settings:

Use this panel to view or change the configuration properties of the selected JMS destination for use with the associated JMS provider.

A JMS destination is used to configure the properties of a JMS destination for the associated generic JMS provider. Connections to the JMS destination are created by the associated JMS connection factory. A JMS destination for use with a generic JMS provider (not the embedded WebSphere JMS provider or WebSphere MQ JMS provider) has the following properties.

To view this administrative console page, click **Resources-> Generic JMS Providers-> provider_name-> JMS Destinations-> destination**

Name:

The name by which the queue is known for administrative purposes within IBM WebSphere Application Server.

Data type String

Type:

Whether this JMS destination is a queue (for point-to-point) or topic (for publish/subscribe).

Select one of the following options:

Queue

A JMS queue destination for point-to-point messaging.

Topic A JMS topic destination for publish/subscribe messaging.

JNDI name:

The JNDI name that is used to bind the queue into the application server's name space.

As a convention, use the fully qualified JNDI name; for example, in the form `jms/Name`, where *Name* is the logical name of the resource.

This name is used to link the platform binding information. The binding associates the resources defined by the deployment descriptor of the module to the actual (physical) resources bound into JNDI by the platform.

Data type String

Description:

A description of the queue, for administrative purposes

Category:

A category used to classify or group this queue, for your IBM WebSphere Application Server administrative records.

Data type String

External JNDI name:

The JNDI name that is used to bind the queue into the application server's name space.

As a convention, use the fully qualified JNDI name; for example, in the form `jms/Name`, where *Name* is the logical name of the resource.

This name is used to link the platform binding information. The binding associates the resources defined by the deployment descriptor of the module to the actual (physical) resources bound into JNDI by the platform.

Data type String

Asynchronous messaging - security considerations

This topic describes considerations that you should be aware of if you want to use security for asynchronous messaging with WebSphere Application Server.

Security for messaging operates as a part of the WebSphere Application Server global security, and is enabled only when global security is enabled.

When global security is enabled, JMS connections made to the JMS provider are authenticated, and access to JMS resources owned by the JMS provider are controlled by access authorizations. Also, all requests to create new connections to the JMS provider must provide a user ID and password for authentication. The user ID and password do not need to be provided by the application. If authentication is successful, then the JMS connection is created; if the authentication fails then the connection request is ended.

Standard J2C authentication is used for a request to create a new connection to the JMS provider. You can specify a Component-managed Authentication Alias and a Container-managed Authentication Alias for each JMS connection factory. The use of the associated J2C authentication data entries depends on the resource authentication (res-auth) setting, as follows:

- If your resource authentication (res-auth) is set to Application, set the alias in the Component-managed Authentication Alias. If the application that tries to create a connection to the JMS provider specifies a user ID and password, those values are used to authenticate the creation request. If the application does not specify a user ID and password, the values defined by the Component-managed Authentication Alias are used. If the connection factory is not configured with a Component-managed Authentication Alias, then you receive a runtime JMS exception when an attempt is made to connect to the JMS provider.
- If your res-auth is set to Container, set the Container-managed Authentication Alias. The values defined by the Container-managed Authentication Alias are used to authenticate the creation request. If you do not specify an alias, then you receive a runtime JMS exception when an attempt is made to connect to the JMS provider.

Restriction:

1. User IDs longer than 12 characters cannot be used for authentication with the embedded WebSphere JMS provider. For example, the default Windows NT user ID, **Administrator**, is not valid for use with embedded WebSphere messaging, because it contains 13 characters. Therefore, an authentication alias for a WebSphere JMS provider connection factory must specify a user ID no longer than 12 characters.
2. If you want to use a WebSphere MQ JMS Provider JMS connection when using Bindings transport mode, you set the property **Transport type=BINDINGS** on the WebSphere MQ Queue Connection Factory. You must also choose one of the following options:
 - To use security credentials, ensure that the user specified is the currently logged on user for the WebSphere Application Server process. If the user specified is not the current logged on user for the WebSphere Application Server process, then the WebSphere MQ JMS Bindings authentication throws the error MQJMS2013 invalid security authentication supplied for MQQueueManager error.
 - Do not specify security credentials. On the WebSphere MQ Connection Factory, ensure that both the **Component-managed Authentication Alias** and the **Container-managed Authentication Alias** properties are not set.

Authorization to access JMS resources owned by the embedded WebSphere JMS provider is controlled by authorization data in the config\integral-jms-authorisations.xml file. For information about editing this file, see Configuring security for the embedded WebSphere JMS provider.

Process choreographer tuning tips

Use the following tuning suggestions to improve performance to the process engine:

- Tune the database using a tuning wizard such as DB2's Performance Wizard.
- Use a faster disk, or use RAID volumes and distribute the MQ logs, MQ data, DB2 logs, DB2 data and WebSphere Application Server transaction logs across separate disks.
- Tune the WebSphere Application Server runtime to efficiently handle data and messages, and use maximum capacity with the following configuration settings:
 - Increase the Java Message Service (JMS) Listener port Max sessions parameter to higher than the default value by Configuring the listener port. Increasing the value enables efficient message handling by the flow message-driven beans. Based on the amount of work and the resources available, a value between 15 and 25 gives best performance and scaling on enterprise systems.
 - Increase the database connection pool size (see "Connection pool settings" in the Information Center) and the Statement cache size (see "Data source settings" in the Information Center) to optimize the data handling. In performance tests, when using a value of 50 for both settings on a NetFinity 5500 machine with an enterprise workload on long running processes, best performance is achieved.

- Increase the JMS provider connection pool size and session pool settings to improve message handling efficiency. In performance tests, both pool values were set to 600 on a NetFinity 5500 machine to increase message handling efficiency.

Designing an enterprise application to use JMS

This topic describes things to consider when designing an enterprise application to use the JMS API directly for asynchronous messaging.

This topic describes things to consider when designing an enterprise application to use the JMS API directly for asynchronous messaging.

1. The application refers to JMS resources that are predefined, as administered objects, to WebSphere Application Server.

Details of JMS resources that are used by enterprise applications are defined to WebSphere Application Server and bound into the JNDI namespace by the WebSphere administrative support. An enterprise application can retrieve these objects from the JNDI namespace and use them without needing to know anything about their implementation. This enables the underlying messaging architecture defined by the JMS resources to be changed without requiring changes to the enterprise application. When designing an enterprise application, you need to identify the details of the following types of JMS resources:

Point-to-Point	Publish/Subscribe
QueueConnectionFactory Queue	TopicConnectionFactory Topic

A connection factory is used to create connections with the JMS provider for a specific JMS queue or topic destination. Each connection factory encapsulates the configuration parameters needed to create a connection to a JMS destination.

For more information about the properties of these JMS resources, see *Configuring JMS provider resources*.

2. The application server pools connections and sessions with the JMS provider to improve performance. This is independent from any WebSphere MQ connection pooling. You need to configure the connection and session pool properties appropriately for your applications, otherwise you may not get the connection and session behavior that you want.
3. Applications can cache JMS connections, sessions, and producers or consumers. Due to the pooling mentioned above this may not give as much of a performance improvement as you might expect. You *must not* cache session handles in stateless session beans that operate in transactions started by a client of the bean. Caching handles in this way causes the bean to be returned to the pool while the session is still involved in the transaction. Also, you should not cache non-durable subscribers due to the restriction mentioned above.
4. A non-durable subscriber can only be used in the same transactional context (for example, a global transaction or an unspecified transaction context) that existed when the subscriber was created. For more information about this context restriction, see *The effect of transaction context on non-durable subscribers*.
5. If you want to use authentication with embedded WebSphere messaging, you cannot have user IDs longer than 12 characters. For example, the default Windows NT user ID, **administrator**, is not valid for use with WebSphere internal messaging, because it contains 13 characters.
6. For messaging operations, you should write application programs that use only references to the interfaces defined in Sun's `javax.jms` package.

JMS defines a generic view of a messaging that maps onto the underlying transport. An enterprise application that uses JMS, makes use of the following interfaces that are defined in Sun's `javax.jms` package:

Connection

Provides access to the underlying transport, and is used to create Sessions.

Session

Provides a context for producing and consuming messages, including the methods used to create MessageProducers and MessageConsumers.

MessageProducer

Used to send messages.

MessageConsumer

Used to receive messages.

The generic JMS interfaces are subclassed into the following more specific versions for Point-to-Point and Publish/Subscribe behavior:

Point-to-Point	Publish/Subscribe
QueueConnection	TopicConnection
QueueSession,	TopicSession,
QueueSender	TopicSender
QueueReceiver	TopicReceiver

The section “J2EE.6.7 Java Message Service (JMS) 1.0 Requirements” of the J2EE specification gives a list of methods that must not be called in Web and EJB containers:

```
javax.jms.Session method setMessageListener
javax.jms.Session method getMessageListener
javax.jms.Session method run
javax.jms.QueueConnection method createConnectionConsumer
javax.jms.TopicConnection method createConnectionConsumer
javax.jms.TopicConnection method createDurableConnectionConsumer
javax.jms.MessageConsumer method getMessageListener
javax.jms.MessageConsumer method setMessageListener
javax.jms.Connection method setExceptionListener
javax.jms.Connection stop
javax.jms.Connection setClientID
```

This method restriction is enforced in IBM WebSphere Application Server by throwing a `javax.jms.IllegalStateException`.

- The following points, as defined in the EJB specification, apply to the use of flags on `createxxxSession` calls:
 - The transacted flag passed on `createxxxSession` is ignored inside a global transaction and all work is performed as part of the transaction. Outside of a transaction the transacted flag is not used and, if set to true, the application should use `session.commit()` and `session.rollback()` to control the completion of the work. In an EJB2.0 module, if the transacted flag is set to true and outside of an XA transaction, then the session is involved in the WebSphere local transaction and the unresolved action attribute of the method applies to the JMS work.
 - Clients cannot use using `Message.acknowledge()` to acknowledge messages. If a value of `CLIENT_ACKNOWLEDGE` is passed on the `createxxxSession` call, then messages are automatically acknowledged by the application server and `Message.acknowledge()` is not used.
- Decide what message selectors are needed.

You can use the JMS message selector mechanism to select a subset of the messages on a queue so that this subset is returned by a receive call. The selector can refer to fields in the JMS message header and fields in the message properties.

- Acting on messages received.

When a message is received, you can act on it as needed by the business logic of the application. Some general JMS actions are to check that the message is of the correct type and extract the content of the message. To extract the content from the body of the message, you need to cast from the generic `Message` class (which is the declared return type of the receive methods) to the more specific subclass, such as `TextMessage`. It is good practice always to test the message class before casting, so that unexpected errors can be handled gracefully.

In this example, the `instanceof` operator is used to check that the message received is of the `TextMessage` type. The message content is then extracted by casting to the `TextMessage` subclass.


```
        if ( inMessage instanceof TextMessage )
```

```
    ...
```

```
        String replyString = ((TextMessage) inMessage).getText();
```

10. Using a listener to receive messages asynchronously.

An alternative to making calls to `QueueReceiver.receive()` is to register a method that is called automatically when a suitable message is available; for example:

```
    ...
    MyClass listener =new MyClass();
    queueReceiver.setMessageListener(listener);
    //application continues with other application-specific behavior.
    ...
```

When a message is available, it is retrieved by the `onMessage()` method on the listener object.

```
import javax.jms.*;
public class MyClass implements MessageListener
{
    public void onMessage(Message message)
    {
        System.out.println("message is "+message);
        //application specific processing here
        ...
    }
}
```

Note: A `MessageListener` can only be used in the client container. (The J2EE specification forbids the use of the JMS `MessageListener` mechanism for the asynchronous receipt of messages in the EJB and Web containers.)

For asynchronous message delivery, the application code cannot catch exceptions raised by failures to receive messages. This is because the application code does not make explicit calls to `receive()` methods. To cope with this situation, you can register an `ExceptionListener`, which is an instance of a class that implements the `onException()` method. When an error occurs, this method is called with the `JMSException` passed as its only parameter.

For more details about using listeners to receive messages asynchronously, see the [Java Message Service Documentation](#).

Note: An alternative to developing your own JMS listener class, you can use a message-driven bean, as described in [Implementing WebSphere enterprise applications that use message-driven beans](#).

11. Warning when receiving messages within a server-side application component.

Take care when performing a JMS `receive()` from a server-side application component if that `receive()` invocation is waiting on a message produced by another application component that is deployed in the same server. Such a JMS `receive()` is synchronous, so blocks until the response message is received.

This type of application design can lead to the consumer/producer problem where the entire set of work threads can be exhausted by the receiving component, which has been blocked waiting for responses, leaving no available worker thread for which to dispatch the application component that would generate the response JMS message.

To illustrate this problem, picture a servlet and a message-driven bean deployed in the same server. When this servlet dispatches a request it sends a message to a queue which is serviced by the message-driven bean (that is, messages produced by the servlet are consumed by the message-driven bean's `onMessage()` method). The servlet subsequently issues a `receive()`, waiting for a reply on a temporary `ReplyTo` queue. The message-driven bean's `onMessage()` method performs a database query and sends back a reply to the servlet on the temporary queue. If a large number of servlet requests occur at once (relative to the number of server worker threads), then it is likely that all available server worker threads will be used to dispatch a servlet request, send a message, and wait for a reply. The application server enters a deadly-embrace condition whereby no

threads remain to process any of the message-driven beans that are now pending. Since the servlets are waiting in blocking receives, the server hangs, likely leading to application failure.

Possible solutions are:

- a. Ensure that the number of worker threads (# of threads per server region * # of server regions per server) exceeds the number of concurrent dispatches of the application component doing the receive() so that there is always a worker thread available to dispatch the message producing component.
- b. Use an application topology that places the receiver application component in a separate server than the producer application component. While worker thread usage can still need to be carefully considered under such a deployment scenario, this separation ensures that there are always threads that cannot be blocked by the message receiving component. There can be other interactions to consider, such as an application server that has multiple applications installed.
- c. Refactor your application to do the message receives from a client component, which will not compete with the producer component for worker threads. Furthermore, the client component can do asynchronous (non-blocking) receives, which are prohibited from J2EE servers. So, for example, the example application above could be refactored to have a client sending messages to a queue and then waiting for a response from the MDB.

The effect of transaction context on non-durable subscribers

A non-durable subscriber can only be used in the same transactional context (for example, a global transaction or an unspecified transaction context) that existed when the subscriber was created. A non-durable subscriber is invalidated whenever a sharing boundary (in general, a local or global transaction boundary) is crossed, resulting in a `javax.jms.IllegalStateException` with message text `Non-durable subscriber invalidated on transaction boundary`.

For example, in the following scenario the non-durable subscriber is invalidated at the begin user transaction. This is because the local transaction context in which the subscriber was created ends when the user transaction begins:

```
...
create subscriber
...
begin user transaction -
...
complete user transaction -
...
use subscriber
...
```

If you want to cache a subscriber (to wait to receive messages that arrived since it was created), then use a durable subscriber (for which this restriction does not apply). Do not cache non-durable subscribers.

Developing a J2EE application to use JMS

Use this task to develop a J2EE application to use the JMS API directly for asynchronous messaging.

This topic gives an overview of the steps needed to develop a J2EE application (servlet or enterprise bean) to use the JMS API directly for asynchronous messaging.

This topic only describes the JMS-related considerations; it does not describe general J2EE application programming, which you should already be familiar with. For detailed information about these steps, and for examples of developing a J2EE application to use JMS, see the *Java Message Service Documentation* and the *WebSphere MQ Using Java* book, SC34-5456.

Details of JMS resources that are used by J2EE applications are defined to WebSphere Application Server and bound into the JNDI namespace by the WebSphere administrative support.

To use JMS, any method of a J2EE application completes the following general steps:

1. Import JMS packages. A J2EE application that uses JMS starts with a number of import statements for JMS, which should include at least the following:

```
import javax.jms.*;           //JMS interfaces
import javax.naming.*;       //Used for JNDI lookup of administered objects
```

2. Get an initial context.

```
try {
    ctx = new InitialContext(env);
    ...
}
```

3. Retrieve administered objects from the JNDI namespace. The `InitialContext.lookup()` method is used to retrieve administered objects (a queue connection factory and the queue destinations); for example, to receive a message from a queue

```
qcf = (QueueConnectionFactory)ctx.lookup( qcfName );
...
inQueue = (Queue)ctx.lookup( qnameIn );
...
```

4. Create a connection to the messaging service provider. The connection provides access to the underlying transport, and is used to create sessions. The `createQueueConnection()` method on the factory object is used to create the connection.

```
connection = qcf.createQueueConnection();
```

The JMS specification defines that connections should be created in the stopped state. Until the connection starts, `MessageConsumers` that are associated with the connection cannot receive any messages. To start the connection, issue the following command:

```
connection.start();
```

5. Create a session, for sending or receiving messages. The session provides a context for producing and consuming messages, including the methods used to create `MessageProducers` and `MessageConsumers`. The `createQueueSession` method is used on the connection to obtain a session. The method takes two parameters:

- A boolean that determines whether or not the session is transacted.
- A parameter that determines the acknowledge mode.

```
boolean transacted = false;
session = connection.createQueueSession( transacted,
                                         Session.AUTO_ACKNOWLEDGE);
```

In this example, the session is not transacted, and it should automatically acknowledge received messages. With these settings, a message is backed out only after a system error or if the application terminates unexpectedly.

The following points, as defined in the EJB specification, apply to these flags:

- The transacted flag passed on `createQueueSession` is ignored inside a global transaction and all work is performed as part of the transaction. Outside of a transaction the transacted flag is not used and, if set to true, the application should use `session.commit()` and `session.rollback()` to control the completion of the work. In an EJB2.0 module, if the transacted flag is set to true and outside of an XA transaction, then the session is involved in the WebSphere local transaction and the unresolved action attribute of the method applies to the JMS work.
- Clients cannot use using `Message.acknowledge()` to acknowledge messages. If a value of `CLIENT_ACKNOWLEDGE` is passed on the `createQueueSession` call, then messages are automatically acknowledged by the application server and `Message.acknowledge()` is not used.

6. Send a message.

- a. Create `MessageProducers` to create messages. For point-to-point messaging the `MessageProducer` is a `QueueSender` that is created by passing an output queue object (retrieved earlier) into the `createSender` method on the session. A `QueueSender` is normally created for a specific queue, so that all messages sent using that sender are sent to the same destination.

```
QueueSender queueSender = session.createSender(inQueue);
```

- b. Create the message. Use the session to create an empty message and add the data passed.

JMS provides several message types, each of which embodies some knowledge of its content. To avoid referencing the vendor-specific class names for the message types, methods are provided on the Session object for message creation.

In this example, a text message is created from the `outString` property:

```
TextMessage outMessage = session.createTextMessage(outString);
```

- c. Send the message.

To send the message, the message is passed to the `send` method on the `QueueSender`:

```
queueSender.send(outMessage);
```

7. Receive replies.

- a. Create a correlation ID to link the message sent with any replies. In this example, the client receives reply messages that are related to the message that it has sent, by using a provider-specific message ID in a `JMSCorrelationID`.

```
messageID = outMessage.getJMSMessageID();
```

The correlation ID is then used in a message selector, to select only messages that have that ID:

```
String selector = "JMSCorrelationID = '"+messageID+"'";
```

- b. Create a `MessageReceiver` to receive messages. For point-to-point the `MessageReceiver` is a `QueueReceiver` that is created by passing an input queue object (retrieved earlier) and the message selector into the `createReceiver` method on the session.

```
QueueReceiver queueReceiver = session.createReceiver(outQueue, selector);
```

- c. Retrieve the reply message. To retrieve a reply message, the `receive` method on the `QueueReceiver` is used:

```
Message inMessage = queueReceiver.receive(2000);
```

The parameter in the `receive` call is a timeout in milliseconds. This parameter defines how long the method should wait if there is no message available immediately. If you omit this parameter, the call blocks indefinitely. If you do not want any delay, use the `receiveNoWait()` method. In this example, the `receive` call returns when the message arrives, or after 2000ms, whichever is sooner.

- d. Act on the message received. When a message is received, you can act on it as needed by the business logic of the client. Some general JMS actions are to check that the message is of the correct type and extract the content of the message. To extract the content from the body of the message, it is necessary to cast from the generic `Message` class (which is the declared return type of the `receive` methods) to the more specific subclass, such as `TextMessage`. It is good practice always to test the message class before casting, so that unexpected errors can be handled gracefully.

In this example, the `instanceof` operator is used to check that the message received is of the `TextMessage` type. The message content is then extracted by casting to the `TextMessage` subclass.

```
if ( inMessage instanceof TextMessage )
```

```
...
```

```
String replyString = ((TextMessage) inMessage).getText();
```

8. Closing down. If the application needs to create many short-lived JMS objects at the Session level or lower, it is important to close all the JMS resources used. To do this, you call the `close()` method on the various classes (`QueueConnection`, `QueueSession`, `QueueSender`, and `QueueReceiver`) when the resources are no longer required.

```
queueReceiver.close();
```

```
...
```

```
queueSender.close();
```

```
...
```

```
session.close();
```

```

    session = null;
...
    connection.close();
    connection = null;

```

9. Publishing and subscribing messages. To use JMS Publish/Subscribe support instead of point-to-point messaging, the general actions are the same; for example, to create a session and connection. The exceptions are that topic resources are used instead of queue resources (such as TopicPublisher instead of QueueSender), as shown in the following example to publish a message:

```

// Creating a TopicPublisher
    TopicPublisher pub = session.createPublisher(topic);
...
    pub.publish(outMessage);
...
    // Closing TopicPublisher
    pub.close();

```

10. Handling errors Any JMS runtime errors are reported by exceptions. The majority of methods in JMS throw JMSEExceptions to indicate errors. It is good programming practice to catch these exceptions and display them on a suitable output.

Unlike normal Java exceptions, a JMSEException can contain another exception embedded in it. The implementation of JMSEException does not include the embedded exception in the output of its toString() method. Therefore, you need to check explicitly for an embedded exception and print it out, as shown in the following example:

```

    catch (JMSEException je)
    {
        System.out.println("JMS failed with "+je);
        Exception le = je.getLinkedException();
        if (le != null)
        {
            System.out.println("linked exception "+le);
        }
    }

```

After you have packaged your application, you can next deploy the application into WebSphere Application Server, as described in Deploying a J2EE application to use JMS.

Developing a JMS client

Use this task to develop a JMS client application to use messages to communicate with enterprise applications.

This topic gives an overview of the steps needed to develop a JMS client application, based on a sample client provided with WebSphere Application Server. This topic only describes the JMS-related considerations; it does not describe general client programming, which you should already be familiar with. For detailed information about these steps, and for examples of developing JMS clients, see the Java Message Service Documentation and the WebSphere MQ *Using Java* book, SC34-5456.

A JMS client assumes that the JMS resources (such as a queue connection factory and queue destination) already exist. A client application can use JMS resources administered by the application server or administered by the client container regardless of whether the client application is running on the same machine as the server or remotely.

- If you want your client application to use server-administered JMS objects, configure the client application to use those resources as Resource Environment References.
- If you want your client application to use client container-administered JMS resources, then configure those resources as Resource References.

For more information about developing client applications and configuring JMS resources for them, see Developing J2EE application client code and related tasks.

To use JMS, a typical JMS client program completes the following general steps:

1. Import JMS packages. An enterprise application that uses JMS starts with a number of import statements for JMS; for example:

```
import javax.naming.Context;
import javax.naming.InitialContext;
import javax.rmi.PortableRemoteObject;
import javax.jms.*;
```

2. Get an initial context.

```
try {
    ctx = new InitialContext(env);
    ...
}
```

3. Define the parameters that the client wants to use; for example, to identify the queue connection factory and to assemble a message to be sent.

```
public class JMSppSampleClient
{
    public static void main(String[] args)
        throws JMSEException, Exception
    {
        String messageID           = null;
        String outString           = null;
        String qcfName             = "java:comp/env/jms/ConnectionFactory";
        String qnameIn             = "java:comp/env/jms/Q1";
        String qnameOut           = "java:comp/env/jms/Q2";
        boolean verbose            = false;

        QueueSession session      = null;
        QueueConnection connection = null;
        Context ctx               = null;

        QueueConnectionFactory qcf = null;
        Queue inQueue             = null;
        Queue outQueue            = null;

        ...
    }
}
```

4. Retrieve administered objects from the JNDI namespace. The `InitialContext.lookup()` method is used to retrieve administered objects (a queue connection factory and the queue destinations):

```
qcf = (QueueConnectionFactory)ctx.lookup( qcfName );
...
inQueue = (Queue)ctx.lookup( qnameIn );
outQueue = (Queue)ctx.lookup( qnameOut );
...
```

5. Create a connection to the messaging service provider. The connection provides access to the underlying transport, and is used to create sessions. The `createQueueConnection()` method on the factory object is used to create the connection.

```
connection = qcf.createQueueConnection();
```

The JMS specification defines that connections should be created in the stopped state. Until the connection starts, `MessageConsumers` that are associated with the connection cannot receive any messages. To start the connection, issue the following command:

```
connection.start();
```

6. Create a session, for sending and receiving messages. The session provides a context for producing and consuming messages, including the methods used to create `MessageProducers` and `MessageConsumers`. The `createQueueSession` method is used on the connection to obtain a session. The method takes two parameters:

- A boolean that determines whether or not the session is transacted.
- A parameter that determines the acknowledge mode.


```

boolean transacted = false;
session = connection.createQueueSession( transacted,
                                         Session.AUTO_ACKNOWLEDGE);

```

In this example, the session is not transacted, and it should automatically acknowledge received messages. With these settings, a message is backed out only after a system error or if the client application terminates unexpectedly.

7. Send the message.

- a. Create MessageProducers to create messages. For point-to-point the MessageProducer is a QueueSender that is created by passing an output queue object (retrieved earlier) into the createSender method on the session. A QueueSender is normally created for a specific queue, so that all messages sent using that sender are sent to the same destination.

```
QueueSender queueSender = session.createSender(inQueue);
```

- b. Create the message. Use the session to create an empty message and add the data passed. JMS provides several message types, each of which embodies some knowledge of its content. To avoid referencing the vendor-specific class names for the message types, methods are provided on the Session object for message creation.

In this example, a text message is created from the outString property, which could be provided as an input parameter on invocation of the client program or constructed in some other way:

```
TextMessage outMessage = session.createTextMessage(outString);
```

- c. Send the message.

To send the message, the message is passed to the send method on the QueueSender:

```
queueSender.send(outMessage);
```

8. Receive replies.

- a. Create a correlation ID to link the message sent with any replies. In this example, the client receives reply messages that are related to the message that it has sent, by using a provider-specific message ID in a JMSCorrelationID.

```
messageID = outMessage.getJMSMessageID();
```

The correlation ID is then used in a message selector, to select only messages that have that ID:

```
String selector = "JMSCorrelationID = '"+messageID+"'";
```

- b. Create a MessageReceiver to receive messages. For point-to-point the MessageReceiver is a QueueReceiver that is created by passing an input queue object (retrieved earlier) and the message selector into the createReceiver method on the session.

```
QueueReceiver queueReceiver = session.createReceiver(outQueue, selector);
```

- c. Retrieve the reply message. To retrieve a reply message, the receive method on the QueueReceiver is used:

```
Message inMessage = queueReceiver.receive(2000);
```

The parameter in the receive call is a timeout in milliseconds. This parameter defines how long the method should wait if there is no message available immediately. If you omit this parameter, the call blocks indefinitely. If you do not want any delay, use the receiveNoWait() method. In this example, the receive call returns when the message arrives, or after 2000ms, whichever is sooner.

- d. Act on the message received. When a message is received, you can act on it as needed by the business logic of the client. Some general JMS actions are to check that the message is of the correct type and extract the content of the message. To extract the content from the body of the message, you need to cast from the generic Message class (which is the declared return type of the receive methods) to the more specific subclass, such as TextMessage. It is good practice always to test the message class before casting, so that unexpected errors can be handled gracefully.

In this example, the instanceof operator is used to check that the message received is of the TextMessage type. The message content is then extracted by casting to the TextMessage subclass.


```
if ( inMessage instanceof TextMessage )
```

```
...  
    String replyString = ((TextMessage) inMessage).getText();
```

9. Closing down. If the application needs to create many short-lived JMS objects at the Session level or lower, it is important to close all the JMS resources used. To do this, you call the close() method on the various classes (QueueConnection, QueueSession, QueueSender, and QueueReceiver) when the resources are no longer required.

```
    queueReceiver.close();  
...  
    queueSender.close();  
...  
    session.close();  
    session = null;  
...  
    connection.close();  
    connection = null;
```

10. Publishing and subscribing messages. To use publish/subscribe support instead of point-to-point messaging, the general client actions are the same; for example, to create a session and connection. The exceptions are that topic resources are used instead of queue resources (such as TopicPublisher instead of QueueSender), as shown in the following example to publish a message:

```
// Creating a TopicPublisher  
    TopicPublisher pub = session.createPublisher(topic);  
...  
    pub.publish(outMessage);  
...  
    // Closing TopicPublisher  
    pub.close();
```

11. Handling errors Any JMS runtime errors are reported by exceptions. The majority of methods in JMS throw JMSEExceptions to indicate errors. It is good programming practice to catch these exceptions and display them on a suitable output.

Unlike normal Java exceptions, a JMSEException can contain another exception embedded in it. The implementation of JMSEException does not include the embedded exception in the output of its toString() method. Therefore, you need to check explicitly for an embedded exception and print it out, as shown in the following example:

```
    catch (JMSEException je)  
    {  
        System.out.println("JMS failed with "+je);  
        Exception le = je.getLinkedException();  
        if (le != null)  
        {  
            System.out.println("linked exception "+le);  
        }  
    }
```

Deploying a J2EE application to use JMS

This topic describes how to deploy a J2EE application to use JMS.

This task description assumes that you have an .EAR file, which contains an application enterprise bean with code for JMS, that can be deployed in WebSphere Application Server.

To deploy a J2EE application to use JMS, complete the following steps:

1. **5.1+** Configure the deployment attributes for the application, as described in Assembling applications with the Assembly Toolkit.
2. Use the WebSphere administrative console to install the application.

This stage is a standard WebSphere Application Server task, as described in Installing applications.

Tuning Java messaging service

You can tune Java messaging service (JMS) run-time components, resources, and the embedded messaging server through the administrative console with the following parameters:

- **Configure the message listening service.**
 - Thread pool
 - Custom properties Application Server Facilities and Non-Application Server Facilities
 - Listener port Maximum sessions
 - Listener port Maximum messages
- **Configure JMS resources.**
 - XA enabled
 - Connection pool size
 - WebSphere MQ queue connection factory Transport type
 - WebSphere MQ topic connection factory Transport type
 - Transaction log directory
 - **WebSphere MQ**
 - **Description:** Includes the WebSphere MQ folder where WebSphere MQ is installed and indirectly where its data files reside since they are always under the directory *WebSphereMQHome/data*. Use this parameter during installation when extra disks are available.
 - **How to view or set:** To view, search your system for a folder named WebSphere MQ. Set during installation.
 - **Default value:** *c:\Program Files\IBM\WebSphere MQ*
 - **Recommended value:** Locate this directory away from any other disk input or output, on the fastest disk available. Redundant Array of Inexpensive Disks (RAID) is recommended.
 - **WebSphere MQ**
 - **Description:** Includes the WebSphere MQ folder where WebSphere MQ is installed and indirectly where its data files reside since they are always under the directory *WebSphereMQHome/data*. Use this parameter during installation when extra disks are available.
 - **How to view or set:** To view, search your system for a folder named WebSphere MQ. Set during installation.
 - **Default value:** *c:\Program Files\IBM\WebSphere MQ*
 - **Recommended value:** Locate this directory away from any other disk input or output, on the fastest disk available. Redundant Array of Inexpensive Disks (RAID) is recommended.
 - **Log buffer pages**
 - **Description:** Controls the size of the queue manager buffer to log file writes in number of pages. Use this parameter when high volumes of messages are being sent through a server.
 - **How to view or set:** In the registry, navigate to **HKEY_LOCAL_MACHINE > SOFTWARE > IBM > MQSeries > CurrentVision > Configuration > QueueManager > QM_name > Log > LogBufferPages**. For embedded JMS, edit the *createmq.properties* setting *LogFilePages* in the *install_root\properties* directory to the desired value. Run the **deletemq** and **createmq** commands to delete or recreate the queue manager. For external JMS, set *LogFilePages* to the desired value. Do this before creating the queue manager in the registry. Run the **amqmdain regsec** command to secure the registry value and create the queue manager.
 - **Default value:** 0
 - **Recommended value:** There is no performance penalty for setting this value too high. Therefore, set the value to its maximum of 512 pages.
 - **Log primary files**
 - **Description:** Controls the number of primary or permanent log files for the queue manager. Use this parameter when high volumes of messages are being sent through a server.
 - **How to view or set:** In the registry, navigate to **HKEY_LOCAL_MACHINE > SOFTWARE > IBM > MQSeries > CurrentVision > Configuration > QueueManager > QM_name > Log > LogPrimaryFiles**. For embedded JMS, edit the *createmq.properties* setting *LogFilePages* in the *install_root\properties* directory to the desired value. Run the **deletemq** and **createmq** commands

to delete or recreate the queue manager. For external JMS, set LogPrimaryFiles to the desired value. Do this before creating the queue manager in the registry. Run the **amqmdain regsec** command to secure the registry value and create the queue manager.

- **Default value:** External: 3, embedded: 0

- **Recommended value:** There is no performance penalty for setting this value too high. Set the value to its maximum of 63 pages.

– Log secondary files

- **Description:** Controls the number of secondary log files for the queue manager. Secondary files are files created when the primary files are not enough and deleted when they are no longer needed.

- **How to view or set:** In the registry, navigate to **HKEY_LOCAL_MACHINE > SOFTWARE > IBM > MQSeries > CurrentVersion > Configuration > QueueManager > QM_name > Log > LogSecondaryFiles**. For embedded JMS, edit the `createmq.properties` setting LogSecondaryFiles in the `install_root\properties` directory to the desired value. Run the **deletemq** and **createmq** commands to delete or recreate the queue manager. For external JMS, set LogPrimaryFiles to the desired value. Do this before creating the queue manager in the registry. Run the **amqmdain regsec** command to secure the registry value and create the queue manager.

- **Default value:** External: 2, embedded: 60

- **Recommended value:** There is a limit of 63 total files from the primary and secondary combined and because secondary logs are slower, setting this 0 is ideal for performance.

- **Log default path**

- **Description:** Controls the location of the queue manager log files. Use this parameter when high volumes of messages are sent through a server.

- **How to view or set:** In the registry, navigate to **HKEY_LOCAL_MACHINE > SOFTWARE > IBM > MQSeries > CurrentVersion > Configuration > QueueManager > QM_name > Log > LogDefaultPath**. For embedded JMS, edit the `createmq.properties` setting LogPath in the `install_root\properties` directory to the desired value. Run the **deletemq** and **createmq** commands to delete or recreate the queue manager. For external JMS, set LogDefaultPath to the desired value. Do this before creating the queue manager in the registry. Run the **amqmdain regsec** command to secure the registry value and create the queue manager.

- **Default value:** `WebSphereMQHome\log`

- **Recommended value:** It is ideal to have a disk dedicated to this task because WebSphere MQ tries to keep the head of the disk positioned at the place in the file where it needs to write next. A fast RAID volume is best.

- **Default queue buffer size**

- **Description:** Controls the size in bytes of an in-memory buffer for nonpersistent queues. Use this parameter when large message sizes are used, or large bursts of messages cause the queue to back up. If the queue backs up past this buffer, messages are flushed out to the disk.

- **How to view or set:** In the registry, navigate to **HKEY_LOCAL_MACHINE > SOFTWARE > IBM > MQSeries > CurrentVersion > Configuration > QueueManager > QM_name > TuningParameters > DefaultQBufferSize**. Embedded JMS is not currently supported; you need MQ5.3 CSD2 and an accompanying Version 5.0 fix. For external JMS, set DefaultQBufferSize to the desired value. Do this before creating the queue manager. Run the **amqmdain regsec** command to secure the registry value and create the queue manager.

- **Default value:** 64K (registry key does not exist)

- **Recommended value:** Set this parameter to accommodate the typical number of messages sitting on the queue at any given time. This should be $\text{numberOfMessages} * (500 + \text{messageSizeInBytes})$. The maximum value is 100MB, but typically 1MB is enough.

- **Default persistent queue buffer size**

- **Description:** Controls the size in bytes of an in-memory buffer for nonpersistent queues. Use this parameter whenever memory is available.

- **How to view or set:** In the registry, navigate to **HKEY_LOCAL_MACHINE > SOFTWARE > IBM > MQSeries > CurrentVersion > Configuration > QueueManager > QM_name > TuningParameters > DefaultPQBufferSize**. Embedded JMS is not currently supported; you need MQ5.3 CSD2 and an accompanying Version 5.0 fix. For external JMS, set

DefaultPQBufferSize to the desired value. Create the queue manager. This is a permanent queue setting. To make a change, delete and recreate the queue.

- **Default value:** 0 (registry key does not exist)
- **Recommended value:** Set this parameter to accommodate the number of typical concurrently processed messages, plus a little more for read-ahead capabilities. Do this by calculating $\text{numOfCocurrentMessages} * (\text{msgSizeInBytes} + 500) * 2$. Typically, 1MB is enough.
- **Maximum channels**
 - **Description:** Controls the allowable number of concurrent CLIENT transport clients. Use this parameter when large numbers of clients are being used.
 - **How to view or set:** In the registry, navigate to **HKEY_LOCAL_MACHINE > SOFTWARE > IBM > MQSeries > CurrentVersion > Configuration > QueueManager > QM_name > Channels > MaxChannels**. Embedded JMS is not currently supported; you need MQ5.3 CSD2 and an accompanying Version 5.0 fix. For external JMS, set MaxChannels to the desired value. Restart the queue manager.
 - **Default value:** External: 100 (registry key does not exist), embedded: 1000
 - **Recommended value:** Set this parameter high enough to contain the maximum number of concurrent JMS clients.
- **Channel application bind type**
 - **Description:** Controls if the channel application is an MQ FASTPATH application. Use this parameter at all times.
 - **How to view or set:** In the registry, navigate to **HKEY_LOCAL_MACHINE > SOFTWARE > IBM > MQSeries > CurrentVersion > Configuration > QueueManager > QM_name > Channels > MQIBindType**. Embedded JMS is not currently supported; you need MQ5.3 CSD2 and an accompanying Version 5.0 fix. For external JMS, set MQIBindType to the desired value. Restart the queue manager.
 - **Default value:** Not FASTPATH (registry key does not exist)
 - **Recommended value:** FASTPATH
- **Configure resources for the embedded WebSphere JMS provider.**
 - **Number of threads**
 - **Description:** With the embedded JMS publications and subscriptions server, this value is the number of threads to use for the publications and subscriptions matching engine, which matches publications to subscribers. Use this parameter when concurrent publications and subscriptions exist that would exceed the capacity of the default value.
 - **How to view or set:**
 1. Open the administrative console.
 2. Click **Servers > Application Servers > server_name**.
 3. Click the *Server Component* > **JMS servers**.
 4. Click **Apply** or **OK**.
 5. Click **Save**.
 6. Stop and restart the application server.
 - **Default value:** 1
 - **Recommended value:** Set this value a little higher than the number of concurrent message publishers. If large numbers of subscribers exist, increasing this value can also provide some benefit.

Troubleshooting WebSphere Messaging

Use this overview task to help resolve a problem that you think is related to the WebSphere Messaging.

To identify and resolve problems that you think are related to WebSphere Messaging, you can use the standard WebSphere Application Server troubleshooting facilities. If you encounter a problem that you think might be related to WebSphere Messaging, complete the following stages. Some problems and their troubleshooting are specific to whether you are using the embedded WebSphere Messaging or WebSphere MQ as the JMS provider.

1. Check for common problems related to WebSphere Messaging. For example, check that the JMS server has been started, that you have added queue names to the list on the JMS server page of the Administrative Console, and that you have successfully installed the WebSphere Messaging function. For tips about solving problems related to the WebSphere Messaging, see *Tips for troubleshooting WebSphere Messaging*. If those tips do not help you fix the problem, complete the following general stages.
2. Check the Release Notes for specific problems and workarounds The section *Possible Problems and Suggested Fixes* of the Release Notes, available from the WebSphere Application Server library web site, is updated regularly to contain information about known defects and their workarounds. Check the latest version of the Release Notes for any information about your problem. If the Release Notes do not contain any information about your problem, you can also search the Technotes database on the WebSphere Application Server web site.
3. Check for WebSphere Messaging error messages.
Check in the application server's SystemOut log at `was_home\logs\server\SystemOut` for error messages with the prefixes MSGS and WMSG.
The associated message reference information provides an explanation and any user actions to resolve the problem. (See "Message reference" in the Information Center.)
4. Check for more informational and error messages that might provide a clue to a related problem. If the JMS server is running, but you have problems accessing JMS resources, check for more error messages and extra details about the problem.
For messages related to WebSphere Messaging, look for the prefixes: MSGS and WMSG.
5. Check your JMS resource configurations If the WebSphere Messaging functions seem to be running properly (the JMS server is running without problems), check that the JMS resources have been configured correctly. For example, check that queue destinations and their connection factories have corresponding JNDI names, that the JNDI names match those configured for the messaging applications, and that the connection factories are configured onto nodes that can provide the JMS resources.
6. Get a detailed exception dump for WebSphere Messaging. If the information obtained in the preceding steps is still inconclusive, you can enable the application server debug trace for the "Messaging" group to provide a detailed exception dump.

Tips for troubleshooting WebSphere Messaging

This topic provides a set of tips to help you troubleshoot problems with WebSphere Messaging.

- The JMS server does not start by default when WebSphere Messaging is installed
- The JMS server tries to start, but fails
- An MDB listener fails to start
- Failure to create a queue connection
- Embedded WebSphere Messaging failed to install on top of an existing WebSphere MQ product
- Problems running JMS applications with security enabled
- Queue manager fails to stop on Redhat Linux
- Application server fails to start in zh_TW.EUC locale on Solaris

The JMS server does not start when starting the WebSphere administrative server

During installation, the system PATH setting is updated to include the WebSphere Messaging directories. If you try to start the WebSphere administrative server from a session that does not use the updated system PATH, the JMS Server fails to start properly. To resolve this after installing WebSphere Application Server stop then restart your host or open a new session that uses the updated system PATH.

Also, when installing WebSphere Application Server, the JMS server is not configured to start automatically by default, except on Windows if the messaging samples are installed. So, generally, if you want to use JMS resources on that host, use the administrative console to configure the JMS server to start automatically. For more information about managing JMS servers, see *Administering WebSphere JMS support*.

Finally, check that the JMS server has started before trying to use WebSphere Messaging. For messages that indicate the JMS server has started successfully, see the following tip The JMS server tries to start, but fails.

The JMS server tries to start, but fails

To see if the JMS Server has started, check the `was_home\logs\server\SystemOut` log. If the JMS server started successfully, you should see the following messages:

```
MSG0050I: Starting the Queue Manager
MSG0051I: Queue Manager open for business
MSG0052I: Starting the Broker
MSG0053I: Broker open for business
```

If the JMS server tries to start, but fails, you should see messages indicating that the queue manager could not start, like the following:

```
MSG0153E: The Queue Manager process could not be started - error: com.ibm.ws.process.exception.
InvalidExecutableException: The system cannot find the file specified.
002: No such file or directory
```

An MDB listener fails to start

If an MDB listener fails to start, you should see the following message:

```
WMSG0019E: Unable to start MDB Listener {0}, JMSDestination {1} : {2}
```

To troubleshoot the cause of an MDB listener not starting, check the following factors:

- Check that the administrative resources have been configured correctly; for example, use the administrative console to check the listener port properties: Destination JNDI name and Connection factory JNDI name. Check that other properties of the listener port, destination, and connection factory are correct.
- Check that the queue exists and has been added to the JMS server.
- Check that the queue manager and JMS server have started.
- Check that the Remote Queue Manager Listener has started.
- If security is enabled, check that a component-managed authentication alias has been specified on the queue connection factory or topic connection factory used by the message-driven bean. This is not required if security is not enabled.
- Check that the user ID used to start the MDB listener has 12 characters or less.

For example, the default Windows NT user ID, **administrator**, is not valid for use with WebSphere embedded messaging, because it contains 13 characters.

(UNIX platforms only) Check that this user ID is a member of the mqm and mqbrkrs groups.

Failure to create a queue connection

If WebSphere Application Server fails to create a queue connection, the SystemOut.log contains error messages like the following:

```
J2CA0046E: Method addNewConnection caught javax.resource.spi.ResourceAdapterInternalException:
createQueueConnection failed
```

Check that the JMS server is running (including that the Internal WebSphere Messaging or WebSphere MQ JMS provider was installed correctly) as described in preceding tips.

Note: In a Network Deployment or Enterprise multi-node cell, the JMS server used by a messaging application can be on a different node to the application. Either check that all JMS servers in the cell have started, or use the Administrative console to determine which JMS server the application is trying to connect to (look at the Node property of the appropriate connection factory), then check that the JMS server has started.

Embedded WebSphere Messaging failed to install on top of an existing WebSphere MQ product

When preparing to install WebSphere Application Server on a host that already has WebSphere MQ installed, you should ensure that WebSphere MQ is at a prerequisite level, as described in the Release Notes and Supported hardware, software, and APIs page of the WebSphere Application Server library web site. You can also check the WebSphere MQ Support service summary Web page at <http://www-3.ibm.com/software/ts/mqseries/support/summary/>.

If you have a problem installing the embedded WebSphere Messaging function, first check the `mq_install.log`. Failures during the WebSphere Messaging prereq stage usually indicate a shortage of space. Failures after this stage are usually due to a conflict between messaging components already installed on the machine and the levels required to support the J2EE 1.3 specification.

If the embedded WebSphere Messaging function installed successfully, you should see messages like the following in `mq_install.log`:

```
...
date time MsiInstallProduct() returning ERROR_SUCCESS (0)
date time ===== Exiting Publish_And Subscribe Install =====
date time <<Function Successful>> return code WASM_ERROR_SUCCESS (0)
date time ===== End of WebSphere Messaging Install Log =====
```

You can also check the `createMQ.log` for any messages that indicate a configuration problem with the installation of WebSphere Messaging.

Problems running JMS applications with security enabled

When trying to run a JMS application with security enabled, you can encounter problems indicated by one of the error messages:

MSG50508E: The JMS Server security service was unable to authenticate userid:

This indicates that the JMS connection has not provided the WebSphere JMS server with any security credentials.

WMSG0019E: Unable to start MDB Listener PSSampleMDB, JMSDestination Sample/JMS/listen : javax.jms.JMSSecurityException:

This indicates that the security credentials supplied are not valid.

In both cases the problem can be removed by doing one of the following:

- If the authentication mechanism is set to `Application`, then the application needs to supply valid credentials.
- If the authentication mechanism is set to `Container`, then you need to configure the JMS `ConnectionFactory` with a container-managed Authentication Alias and ensure that the associated username and password are valid.

MQJMS2013 invalid security authentication supplied for MQQueueManager

If using a WebSphere MQ JMS Provider JMS connection when using Bindings transport mode, and the user specified is not the current logged on user for the WebSphere Application Server process, then the WebSphere MQ JMS Bindings authentication throws the error MQJMS2013 invalid security authentication supplied for MQQueueManager.

If you want to use a WebSphere MQ JMS Provider JMS connection when using Bindings transport mode, you set the property **Transport type=BINDINGS** on the WebSphere MQ Queue Connection Factory. You must also choose one of the following options:

- To use security credentials, ensure that the user specified is the currently logged on user for the WebSphere Application Server process.
- Do not specify security credentials. On the WebSphere MQ Connection Factory, ensure that both the **Component-managed Authentication Alias** and the **Container-managed Authentication Alias** properties are not set.

For more information about messaging security, see [Asynchronous messaging - security considerations](#).

Queue manager fails to stop on Redhat Linux

When trying to stop an application server on Redhat Linux, the queue manager can hang with a Java core dump, and the last message in the SystemOut.log file is Stopping Queue manager....

This is caused by a known RedHat problem (https://bugzilla.linux.ibm.com/show_bug.cgi?id=2336), that was introduced in libstdc++-2.96-116.7.2 and beyond.

The workaround is to go back to the libstdc++-2.96-108.1 level.

Application server fails to start in zh_TW.EUC locale on Solaris

If you have set the locale to zh_TW.EUC on Solaris, and are using the WebSphere embedded JMS provider or WebSphere MQ as the JMS provider, you can encounter problems that stop application servers starting up.

If you intend using the WebSphere embedded JMS provider or WebSphere MQ as the JMS provider on Solaris, do not set the LANG and LC_ALL variables to zh_TW.EUC (Traditional Chinese locale) to avoid problems when starting application servers. Set the LANG and LC_ALL variables to zh_TW instead of zh_TW.EUC.

Using message-driven beans in applications

WebSphere Application Server supports asynchronous messaging as a method of communication based on the Java Message Service (JMS) programming interface.

Message-driven beans (a type of enterprise bean defined in the EJB 2.0 specification) extend the base JMS support and the Enterprise JavaBean component model to provide automatic asynchronous messaging. When a message arrives on a destination, a listener passes the message to a new instance of a user-developed message-driven bean for processing.

You can use WebSphere Studio Application Developer to develop applications that use message-driven beans. You can use the WebSphere Application Server runtime tools, like the administrative console, to deploy and administer applications that use message-driven beans.

For more information about implementing WebSphere enterprise applications that use message-driven beans, see the following topics:

- An overview of message-driven beans
- Designing an enterprise application to use a message-driven bean
- Developing an enterprise application to use a message-driven bean
- Deploying an enterprise application to use a message-driven bean
- Configuring message listener resources for message-driven beans
- Troubleshooting problems with message-driven beans

Message-driven beans - an overview

WebSphere Application Server supports automatic asynchronous messaging with message-driven beans (a type of enterprise bean defined in the EJB 2.0 specification). Messaging with message-driven beans is shown in the figure “Message-driven beans - an overview.”

The support for message-driven beans is based on the message listener service, which comprises a listener manager that controls and monitors one or more listeners. Each listener monitors a JMS destination for incoming messages. When a message arrives on the destination, the listener passes the

message to a new instance of a user-developed message-driven bean (an enterprise bean) for processing. The listener then looks for the next message without waiting for the bean to return.

Messages arriving at a destination being processed by a listener have no client credentials associated with them; the messages are anonymous. Security depends on the role specified by the RunAs Identity for the message-driven bean as an EJB component. For more information about EJB security, see "Enterprise bean component security" in the Information Center.

You are recommended to develop a message-driven bean to delegate the business processing of incoming messages to another enterprise bean, to provide clear separation of message handling and business processing. This also enables the business processing to be invoked by either the arrival of incoming messages or, for example, from a WebSphere J2EE client.

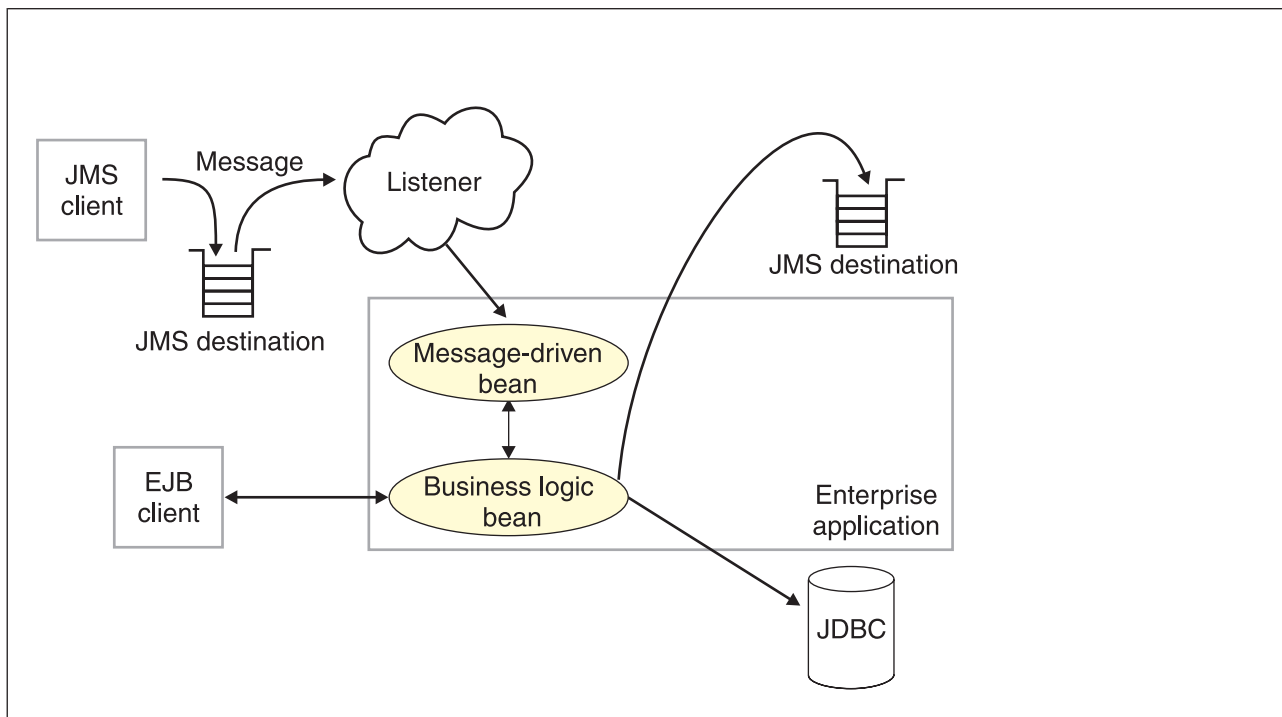


Figure 8. Message-driven beans and the message listener service. This figure shows an incoming message being passed by a JMS listener to a message-driven bean, which passes the message on to a business logic bean for business processing. This messaging is controlled by the listener manager. For more information, see the text that accompanies this figure.

Message-driven beans - components

The WebSphere Application Server support for message-driven beans is based on JMS message listeners and the message listener service, and builds on the base support for JMS. The main components of WebSphere Application Server support for message-driven beans are shown in the following figure and described after the figure:

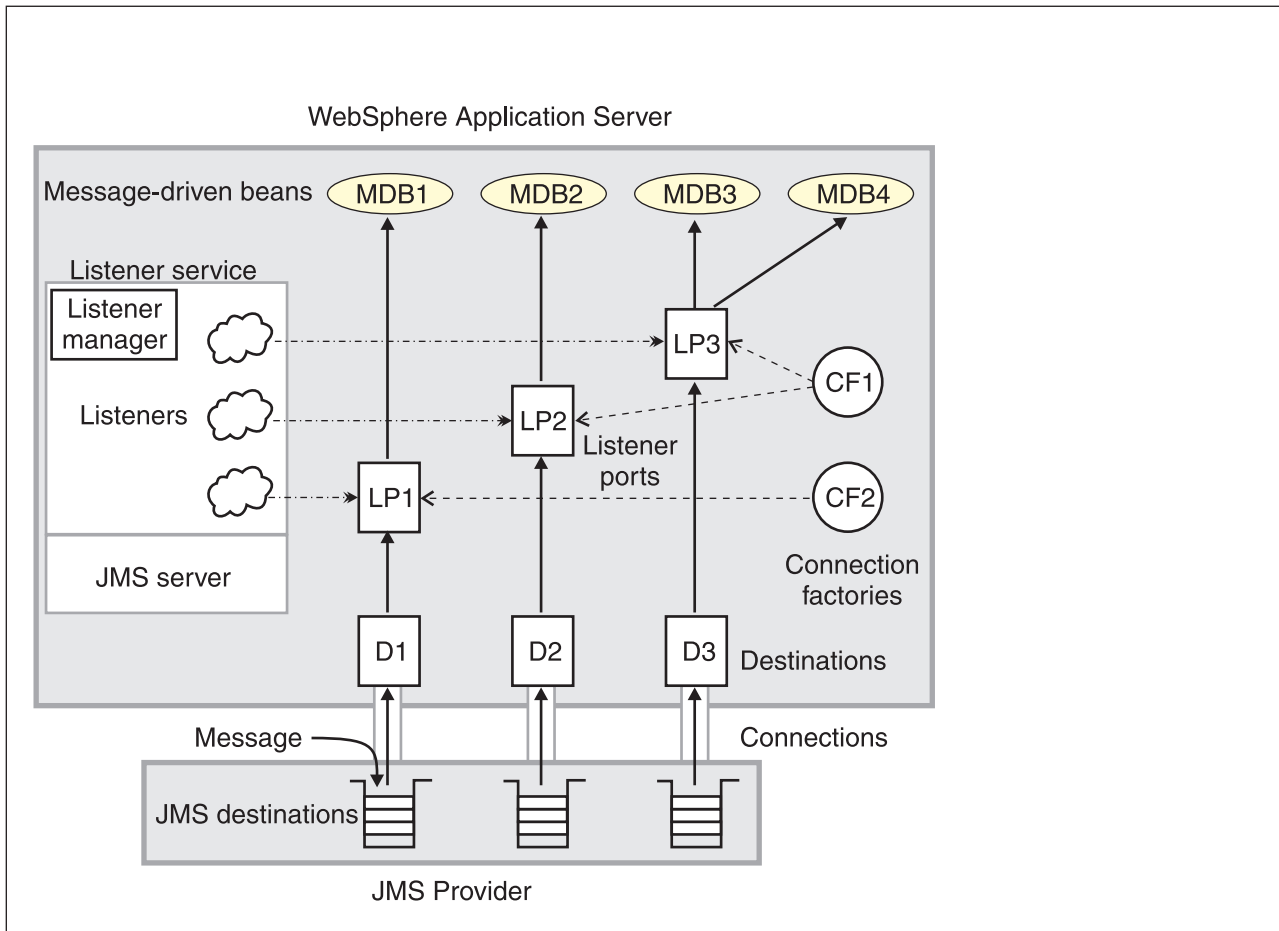


Figure 9. The main components for message-driven beans. This figure shows the main components of WebSphere support for message-driven beans, from JMS provider through a connection to a destination, listener port, then deployed message-driven bean that processes the message retrieved from the destination. Each listener port defines the association between a connection factory, destination, and a deployed message-driven bean. The other main components are the message listener service, which comprises a listener for each listener port, all controlled by the same listener manager. For more information, see the text that accompanies this figure.

The message listener service is an extension to the JMS functions of the JMS provider and provides a listener manager, which controls and monitors one or more JMS listeners.

Each listener monitors either a JMS queue destination (for point-to-point messaging) or a JMS topic destination (for publish/subscribe messaging).

A connection factory is used to create connections with the JMS provider for a specific JMS queue or topic destination. Each connection factory encapsulates the configuration parameters needed to create a connection to a JMS destination.

A listener port defines the association between a connection factory, a destination, and a deployed message-driven bean. Listener ports are used to simplify the administration of the associations between these resources.

When a deployed message-driven bean is installed, it is associated with a listener port and the listener for a destination. When a message arrives on the destination, the listener passes the message to a new instance of a message-driven bean for processing.

When an application server is started, it initializes the listener manager based on the configuration data. The listener manager creates a dynamic session thread pool for use by listeners, creates and starts

listeners, and during server termination controls the cleanup of listener message service resources. Each listener completes several steps for the JMS destination that it is to monitor, including:

- Creating a JMS server session pool, and allocating JMS server sessions and session threads for incoming messages.
- Interfacing with JMS ASF to create JMS connection consumers to listen for incoming messages.
- If specified, starting a transaction and requesting that it is committed (or rolled back) when the EJB method has completed.
- Processing incoming messages by invoking the `onMessage()` method of the specified enterprise bean.

Message-driven beans - transaction support

Message-driven beans can handle messages read from JMS destinations within the scope of a transaction. If transaction handling is specified for a JMS destination, the JMS listener starts a global transaction *before* it reads any incoming message from that destination. When the message-driven bean processing has finished, the JMS listener commits or rolls back the transaction (using JTA transaction control).

Note:

- All messages retrieved from a specific destination have the same transactional behavior.

If messages are queued to be sent within a global transaction they are sent when the transaction is committed. If the processing of a message causes the transaction to be rolled back, then the message that caused the bean instance to be invoked is left on the JMS destination.

You can configure the **Maximum retries** property of the listener port to define the maximum number of times the listener attempts to read a message from a destination. When the Max retries limit is reached, the listener for that destination is stopped. When you have resolved the problem, you must then restart the listener.

Designing an enterprise application to use message-driven beans

This topic describes things to consider when designing an enterprise application to use message-driven beans.

The considerations in this topic are based on a generic enterprise application that uses one message-driven bean to retrieve messages from a JMS queue destination and passes the messages on to another enterprise bean that implements the business logic.

To design an enterprise application to use message-driven beans, complete the following steps:

1. Identify the JMS resources that the application is to use. This helps to identify the properties of resources that need to be used within the application and configured as application deployment descriptors or within WebSphere Application Server.

JMS resource type	Properties
Queue connection factory	Name: SamplePtoPQueueConnectionFactory JNDI Name: Sample/JMS/QCF
Queue destination	Name: Q1 JNDI Name: Sample/JMS/Q1
Listener port (for the destination)	Name: SamplePtoPListenerPort Connection Factory JNDI Name: Sample/JMS/QCF Destination JNDI Name: Sample/JMS/Q1 Maximum Sessions: 5 Maximum Retries: 10 Maximum Messages: 1

JMS resource type	Properties
Message-driven bean (deployment properties)	Name: JMSppSampleMDBBean Transaction type: Container Destination type: Queue Listener port name: SamplePtoPListenerPort
Business logic bean	Name: MyLogicBean

Ensure that you use consistent values where needed; for example, the JNDI names for the connection factory and destination must be the same for both those resources and the equivalent properties of the listener port.

2. Separation of business logic. You are recommended to develop a message-driven bean to delegate the business processing of incoming messages to another enterprise bean. This provides clear separation of message handling and business processing. This also enables the business processing to be invoked by either the arrival of incoming messages or, for example, from a WebSphere J2EE client.

3. Security considerations.

Messages arriving at a destination being processed by a listener have no client credentials associated with them; the messages are anonymous. Security depends on the role specified by the RunAs Identity for the message-driven bean as an EJB component. For more information about EJB security, see "Enterprise bean component security" in the Information center.

4. General JMS considerations For Publish/Subscribe messaging, choose the JMS server port to be used depending on your needs for transactions or performance:

Queued port

The TCP/IP port number of the listener port used for all point-to-point and Publish/Subscribe support.

Direct port

The TCP/IP port number of the listener port used for direct TCP/IP connection (non-transactional, non-persistent, and non-durable subscriptions only) for Publish/Subscribe support.

Note: Message-driven beans cannot use the direct listener port for Publish/Subscribe support. Therefore, any topic connection factory configured with **Portset** to **Direct** cannot be used with message-driven beans.

A non-durable subscriber can only be used in the same transactional context (for example, a global transaction or an unspecified transaction context) that existed when the subscriber was created. For more information about this context restriction, see The effect of transaction context on non-durable subscribers.

Developing an enterprise application to use message-driven beans

Use this task to develop an enterprise application to use a message-driven bean. The message-driven bean is invoked by a JMS listener when a message arrives on the input queue that the listener is monitoring.

You are recommended to develop the message-driven bean to delegate the business processing of incoming messages to another enterprise bean, to provide clear separation of message handling and business processing. This also enables the business processing to be invoked by either the arrival of incoming messages or, for example, from a WebSphere J2EE client. Responses can be handled by another enterprise bean acting as a sender bean, or handled in the message-driven bean.

You develop an enterprise application to use a message-driven bean like any other enterprise bean, except that a message-driven bean does not have a home interface or a remote interface.

This topic describes how to develop a completely new message-driven bean class. If you have a WAS 4.0 enterprise application that uses the JMS listener, you can migrate that application to use message-driven beans, as described in *Migrating a WAS 4.0 JMS listener application to use message-driven beans*.

For more information about writing the message-driven bean class, see *Creating a message-driven bean* in the WebSphere Studio help bookshelf.

To develop an enterprise application to use a message-driven bean, complete the following steps:

1. Creating the Enterprise Application project, as described in the WebSphere Studio article *Creating an Enterprise Application project*.
2. Creating the message-driven bean class.

You can use the New Enterprise Bean wizard of WebSphere Studio Application Developer to create an enterprise bean with a bean type of Message-driven bean. The wizard creates appropriate methods for the type of bean.

By convention, the message bean class is named *nameBean*, where *name* is the name you assign to the message bean; for example:

```
public class MyJMSPMDBBean implements MessageDrivenBean, MessageListener
```

The message-driven bean class must define and implement the following methods:

- `onMessage(message)`, which must meet the following requirements:
 - The method must have a single argument of type `javax.jms.Message`.
 - The throws clause must *not* define any application exceptions.
 - If the message-driven bean is configured to use bean-managed transactions, it must call the `javax.transaction.UserTransaction` interface to scope the transactions. Because these calls occur inside the `onMessage()` method, the transaction scope does not include the initial message receipt. This means the application server is given one attempt to process the message.

To handle the message within the `onMessage()` method (for example, to pass the message on to another enterprise bean), you use standard JMS. (This is known as bean-managed messaging.)

- `ejbCreate()`

You must define and implement an `ejbCreate` method for each way in which you want a new instance of an enterprise bean to be created.

- `ejbRemove()`.

This method is invoked by the container when a client invokes the remove method inherited by the enterprise bean's home interface from the `javax.ejb.EJBHome` interface. This method must contain any code that you want to execute before an enterprise bean instance is removed from the container (and the associated data is removed from the data source).

For example, the following code extract shows how to access the text and the JMS MessageID, from a JMS message of type `TextMessage`:

```

public void onMessage(javax.jms.Message msg)
{
    String text      = null;
    String messageID = null;

    try
    {
        text = ((TextMessage)msg).getText();

        System.out.println("senderBean.onMessage(), msg text2: "+text);

        //
        // store the message id to use as the Correlator value
        //
        messageID = msg.getJMSMessageID();

        // Call a private method to put the message onto another queue
        putMessage(messageID, text);
    }
    catch (Exception err)
    {
        err.printStackTrace();
    }
    return;
}

```

Figure 10. Code example: The `onMessage()` method of a message bean. This figure shows a code extract for a basic `onMessage()` method of a sample message-driven bean. The method unpacks the incoming text message to extract the text and message identifier and calls a private `putMessage` method (defined within the same message bean class) to put the message onto another queue.

The result of this step is a message-driven bean that can be assembled into an .EAR file for deployment.

3. Assembling and packaging the application for deployment.

You can use WebSphere Studio to assemble and package the application for deployment.

The result of this task is an .EAR file, containing an application message-driven bean, that can be deployed in WebSphere Application Server.

After you have developed an enterprise application to use message-driven beans, configure and deploy the application; for example, define the listener ports for the message-driven beans and, optionally, change the deployment descriptor attributes for the application. For more information about configuring and deploying an application that uses message-driven beans, see [Deploying an enterprise application to use message-driven beans](#)

Migrating a JMS listener application to use message-driven beans

Use this task to migrate an enterprise application that uses message beans with the JMS Listener from WebSphere Application Server 4.0 to use EJB 2.0 message-driven beans.

This task uses a command line utility, `mb2mdb`, that takes as its input either a deployed `MessageBean.jar` module or a deployed Enterprise Application (.ear) that contains a message bean, along with the JMS listener configuration XML file that defines the WebSphere Application Server 4.0 message beans. The result is a new .jar/.ear module that can then be deployed directly into a WebSphere Application Server 5.0 application server.

You can display the usage help for the migration utility, by typing the command `mb2mdb` at a command line.

To migrate a WebSphere Application Server 4.0 enterprise application that uses message beans to use EJB 2.0 message-driven beans, type the following command at an operating system command line:

```
mb2mdb inputMB.jar-ear jmsListenerConfig.xml workingDirectory outputMDB.jar-ear options
```


Where:

inputMB.jar-ear

The name of the deployed WebSphere Application Server 4.0 jar or ear file containing a stateless session message bean.

jmsListenerConfig.xml

The name of the XML configuration file used to configure the WebSphere Application Server 4.0 JMS listeners.

workingDirectory

The name of a new or existing directory that is used to generate the new message-driven bean and package the outputMDB.jar or .ear file.

Note: By default, the tool clears the working directory after it has completed. If you want to preserve the contents of the working directory, you must specify the `-keep` option.

outputMDB.jar-ear

The name of the output .jar or .ear file for the migrated message-driven bean application.

options

An optional set of parameters that you can use to control the mb2mdb utility.

-keep This prevents the tool from clearing out the working directory after completion.

-verbose

This causes the tool to display informational messages as to the progress of the migration and its parameters.

-map listenerHome=bindingHome

This option provides a mechanism to map between the JNDIHomeName specified for a listener in the JMS listener configuration XML file and the default binding home name specified in the *inputMB.jar-ear* file.

If the *jmsListenerConfig.xml* file contains a deployed EJB home JNDI name that is different to the default binding within the *inputMB.jar-ear*, use this option to map between the two names.

This enables you to install the output .jar or .ear file for the message-driven bean into an application server and bind the bean with a different JNDIHomeName than is specified in the bean's bindings.xmi.

The result of this task is a new .jar or .ear file for a message-driven bean that can then be deployed directly into a WebSphere Application Server 5.0 application server.

To successfully install the .jar or .ear file, you need to bind the message-driven bean against a listener port defined to the message listener service of the application server. You need to have used the WebSphere Application Server administrative console to define the listener port, which defines the JMS connection factory and destination that a message-driven bean bound to it listens on. For more information about installing and configuring a .jar or .ear file for a message-driven bean, see *Deploying an enterprise application to use message-driven beans*.

Deploying an enterprise application to use message-driven beans

Use this task to deploy an enterprise application to use message-driven beans.

This task description assumes that you have an .EAR file, which contains an application enterprise bean with code for message-driven beans, that can be deployed in WebSphere Application Server.

To deploy an enterprise application to use message-driven beans, complete the following steps:

1. Use the WebSphere administrative console to define the listener ports for the application, as described in *Adding a new listener port*.

2. **5.1+** For each message-driven bean in the application, configure the deployment attributes to match the listener port definitions, as described in *Configuring deployment attributes using the Assembly Toolkit*.

3. Use the WebSphere administrative console to install the application.

This stage is a standard WebSphere Application Server task, as described in *Installing a new application*.

When you install the application, you are prompted to specify the name of the listener port that the application is to use for late responses. Select the listener port, then click **OK**.

Configuring deployment attributes using the Assembly Toolkit

Use this task to configure the message-driven beans deployment attributes for an enterprise bean, to override the deployment attributes defined within the application EAR file.

You can configure the deployment attributes of an application by using the Deployment Descriptor Editor of WebSphere Studio Application Developer or the Assembly Toolkit.

This topic describes the use of the Assembly Toolkit to configure the deployment attributes of an application. This task description assumes that you have an EAR file, which contains an application enterprise bean developed as a message-driven bean, that can be deployed in WebSphere Application Server. For more details about using the Assembly Toolkit, see *Assembling applications with the Assembly Toolkit*.

To configure the message-driven beans deployment attributes for an enterprise bean, use the Assembly Toolkit to configure the deployment attributes of the application to match the listener port definitions:

1. Start the Assembly Toolkit.
2. Create or edit the application EAR file. For example, to change attributes of an existing application, use the import wizard to import the EAR file into the Assembly Toolkit. To start the import wizard:
 - a. Click **File-> Import-> EAR file**
 - b. Click **Next**, then select the EAR file.
 - c. Click **Finish**
3. In the J2EE Hierarchy view, right-click the EJB module for the message-driven bean, then click **Open With > Deployment Descriptor Editor**. A property dialog notebook for the message-driven bean is displayed in the property pane.
4. Specify general deployment properties.
 - a. In the property pane, select the Beans tab.
 - b. Specify the following properties:

Transaction type

Whether the message bean manages its own transactions or the container manages transactions on behalf of the bean. All messages retrieved from a specific destination have the same transactional behavior. To enable the transactional behavior that you want, you must configure the JMS destination with the same transactional behavior as you configure for the message bean.

Bean The message bean manages its own transactions

Container

The container manages transactions on behalf of the bean

5. Specify advanced deployment properties.
 - a. Specify the following properties:

Message selector

The JMS message selector to be used to determine which messages the message bean receives; for example:

```
JMSType='car' AND color='blue' AND weight>2500
```

The selector string can refer to fields in the JMS message header and fields in the message properties. Message selectors cannot reference message body values.

Acknowledge mode

How the session acknowledges any messages it receives.

This property applies only to message-driven beans that uses bean-managed transaction demarcation (**Transaction type** is set to Bean).

Auto Acknowledge

The session automatically acknowledges a message when it has either successfully returned from a call to receive, or the message listener it has called to process the message successfully returns.

Dups OK Acknowledge

The session lazily acknowledges the delivery of messages. This is likely to result in the delivery of some duplicate messages if JMS fails, so it should be used only by consumers that are tolerant of duplicate messages.

As defined in the EJB specification, clients cannot use using Message.acknowledge() to acknowledge messages. If a value of CLIENT_ACKNOWLEDGE is passed on the createxxxSession call, then messages are automatically acknowledged by the application server and Message.acknowledge() is not used.

Destination type

Whether the message bean uses a queue or topic destination.

Queue

The message bean uses a queue destination.

Topic The message bean uses a topic destination.

Subscription durability

Whether a JMS topic subscription is durable or non-durable.

Durable

A subscriber registers a durable subscription with a unique identity that is retained by JMS. Subsequent subscriber objects with the same identity resume the subscription in the state it was left in by the earlier subscriber. If there is no active subscriber for a durable subscription, JMS retains the subscription's messages until they are received by the subscription or until they expire.

Nondurable

Non-durable subscriptions last for the lifetime of their subscriber object. This means that a client sees the messages published on a topic only while its subscriber is active. If the subscriber is not active, the client is missing messages published on its topic.

A non-durable subscriber can only be used in the same transactional context (for example, a global transaction or an unspecified transaction context) that existed when the subscriber was created. For more information about this context restriction, see The effect of transaction context on non-durable subscribers.

6. Specify bindings deployment properties.
 - a. Specify the following property:
Listener port name
Type the name of the listener port for this message-driven bean.
7. Save your changes to the deployment descriptor.
 - a. Close the deployment descriptor editor.
 - b. When prompted, click **Yes** to indicate that you want to save changes to the deployment descriptor.
8. Verify the archive files.
9. Generate code for deployment for EJB modules or for enterprise applications that use EJB modules.
10. Optional: Test your completed module on a WebSphere Application Server installation. Right-click a module, click **Run on Server**, and follow the instructions in the displayed wizard. Note that **Run on**

Server works on the Windows, Linux/Intel, and AIX operating systems only; you cannot deploy remotely from the Assembly Toolkit to a WebSphere Application Server installation on a UNIX operating system such as Solaris.

Important

Important: Use **Run On Server** for unit testing only. Assembly Server Toolkit controls the WebSphere Application Server installation and, when an application is published remotely, the Toolkit overwrites the server configuration file for that server. Do not use on production servers.

For instructions on remote testing, see the article “Setting Up a Remote WebSphere Application Server in WebSphere Studio V5” at http://www7b.boulder.ibm.com/wsdd/techjournal/0303_yuen/yuen.html.

After assembling your application, use a systems management tool to deploy the EAR file onto the application server that is to run the application; for example, using the administrative console as described in Deploying and managing applications.

Configuring message listener resources for message-driven beans

Use the following tasks to configure resources needed by the message listener service to support message-driven beans.

- Configuring the message listener service
- Adding a new listener port
- Configuring a listener port
- Configuring security for message-driven beans

Configuring the message listener service

Use this task to configure the properties of the message listener service for an application server.

To configure the properties of the message listener service for an application server, use the administrative console to complete the following steps:

1. In the navigation pane, select **Servers-> Application Servers** This displays a table of the application servers in the administrative domain.
2. In the content pane, click the name of the application server. This displays the properties of the application server in the content pane.
3. In the Additional Properties table, select **Message Listener Service** This displays the Message Listener Service properties in the content pane.
4. Specify appropriate properties of the message listener service.
5. Optional: Specify any of the following optional properties that you need, as **Custom properties** of the message listener service: `NON.ASF.RECEIVE.TIMEOUT`, `MQJMS.POOLING.TIMEOUT`, `MQJMS.POOLING.THRESHOLD`, `MAX.RECOVERY.RETRIES`, and `RECOVERY.RETRY.INTERVAL`.
For more information about these custom properties, see Custom Properties.
6. Click **OK**.
7. Save your configuration.
8. To have the changed configuration take effect, stop then restart the Application Server.

Message listener service:

The message listener service is an extension to the JMS functions of the JMS provider. It provides a listener manager that controls and monitors one or more JMS listeners, which each monitor a JMS destination on behalf of a deployed message-driven bean.

This panel displays links to the Additional Properties pages for Listener Ports, Thread Pool, and Custom Properties for the message listener service.

To view this administrative console page, click **Servers-> Application Servers-> *application_server*-> Message Listener Service**

Thread pool:

Controls the maximum number of threads the Message Listener Service is allowed to run. Select this link to display the service thread pool properties.

Adjust this parameter when multiple message-driven beans are deployed in the same application server and the sum of their maximum session values exceeds the default value of 10.

Data type	Integer
Units	Not applicable
Default	Minimum: 10, maximum 50
Range	Not applicable
Recommended	Set the minimum to the sum of all message-driven beans maximum session values. Set the maximum to anything equal or greater than the minimum.

Custom Properties:

An optional set of name and value pairs for custom properties of the message listener service.

You can use the Custom properties page to define the following properties for use by the message listener service.

- NON.ASF.RECEIVE.TIMEOUT
- MQJMS.POOLING.TIMEOUT
- MQJMS.POOLING.THRESHOLD
- MAX.RECOVERY.RETRIES
- RECOVERY.RETRY.INTERVAL

Message listener service custom properties:

Use this panel to view or change an optional set of name and value pairs for custom properties of the message listener service.

To view this administrative console page, click **Servers-> *application_server*-> Message Listener Service-> (In content pane, under Additional Properties) Custom Properties**

You can use the Custom properties page to define the following properties for use by the message listener service.

- NON.ASF.RECEIVE.TIMEOUT
- MQJMS.POOLING.TIMEOUT
- MQJMS.POOLING.THRESHOLD
- MAX.RECOVERY.RETRIES
- RECOVERY.RETRY.INTERVAL

NON.ASF.RECEIVE.TIMEOUT:

The timeout in milliseconds for synchronous message receives performed by message-driven bean listener sessions in the non-ASF mode of operation.

You should set this property to a non-zero value only if you want to enable the non-ASF mode of operation for all message-driven bean listeners on the application server.

The message listener service has two modes of operation, Application Server Facilities (ASF) and non-Application Server Facilities (non-ASF).

- The ASF mode is meant to provide concurrency and transactional support for applications. For publish/subscribe message-driven beans, the ASF mode provides better throughput and concurrency, because in the non-ASF mode the listener is single-threaded.
- The non-ASF mode is mainly for use with generic JMS providers that do not support JMS ASF, which is an optional extension to the JMS specification. The non-ASF mode is also transactional but, because the path length is shorter than the ASF mode, usually provides improved performance.

Use non-ASF if:

- Your generic JMS provider does not provide JMS ASF support
- You are using message-driven beans with WebSphere topic connections with the DIRECT port, because the embedded publish/subscribe broker using that port does not support XA transactions or JMS ASF.
- Message order is a strict requirement

Data type	Integer
Units	Milliseconds
Default	ASF mode (custom property not created)
Range	0 or greater milliseconds
	0 non-ASF mode is disabled
	1 or more
	The timeout in milliseconds for non-ASF message-driven bean listener synchronous session receives
Recommended	If a transaction timeout occurs, the message must recycle causing extra work. If you want to use the non-ASF mode, set this property to lower than the transaction timeout, but leave spare at least the maximum duration of your message-driven bean's onMessage() method. For example, if your message-driven bean's onMessage() method typically takes a maximum of 10 seconds, and the transaction timeout is set to 120 seconds, you might set the NON.ASF.RECEIVE.TIMEOUT property to no more than 110000 (110000 milliseconds, that is 110 seconds).

MQJMS.POOLING.TIMEOUT:

The number of milliseconds after which a connection in the pool is destroyed if it has not been used.

An MQSimpleConnectionManager allocates connections on a most-recently-used basis, and destroys connections on a least-recently-used basis. By default, a connection is destroyed if it has not been used for five minutes.

Data type	Integer
Units	Milliseconds
Default	5 minutes
Range	

MQJMS.POOLING.THRESHOLD:

The maximum number of unused connections in the pool.

An MQSimpleConnectionManager allocates connections on a most-recently-used basis, and destroys connections on a least-recently-used basis. By default, a connection is destroyed if there are more than ten unused connections in the pool.

Data type	Integer
Units	Number of connections
Default	10
Range	

MAX.RECOVERY.RETRIES:

The maximum number of times that the listener service tries to get a message from a listener port before the associated listener is stopped, in the range 0 through 2147483647.

Data type	Integer
Units	Retry attempts
Default	0 (no retries)
Range	0 (no retries) through 2147483647

RECOVERY.RETRY.INTERVAL:

The time in seconds between retry attempts by the listener service to get a message from a listener port.

Data type	Integer
Units	Seconds
Default	10
Range	1 through 2147483647

Message listener port collection:

The message listener ports configured in the administrative domain

This panel displays a list of the message listener ports configured in the administrative domain. Each listener port is used with a message-driven bean to automatically receive messages from an associated JMS destination. You can use this panel to add new listener ports or to change the properties of existing listener ports. For more information about the property fields for listener ports, see Listener port properties.

To view this administrative console page, click **Servers-> application_server-> Message Listener Service-> Listener Ports**

Listener port settings:

A listener port is used to simplify administration of the association between a connection factory, destination, and deployed message-driven bean.

Use this panel to view or change the configuration properties of the selected listener port.

To view this administrative console page, click **Servers-> Application Servers-> application_server-> Message Listener Service-> Listener Ports-> listener_port**

Name:

The name by which the listener port is known for administrative purposes.

Data type	String
------------------	--------

Default Null

Initial state:

The state that you want the listener port to have when the application server is next restarted

Data type	Enum
Units	Not applicable
Default	Started
Range	Started When the application server is next started, the listener port is started automatically. Stopped When the application server is next started, the listener port is not started automatically. If message-driven beans are to use this listener port on the application server, the system administrator must start the port manually or select the Started value of this property then restart the application server.

Description:

A description of the listener port, for administrative purposes within IBM WebSphere Application Server.

Data type	String
Default	Null

Connection factory JNDI name:

The JNDI name for the JMS connection factory to be used by the listener port; for example, jms/connFactory1.

Data type	String
Default	Null

Destination JNDI name:

The JNDI name for the destination to be used by the listener port; for example, jms/destn1.

If the extended messaging service is to use this listener port to handle late responses, the value of this property must match the JMS response destination on the output port used by the sender bean.

You cannot use a temporary destination for late responses.

Data type	String
Default	Null

Maximum sessions:

Specifies the maximum number of concurrent sessions that a listener can have with the JMS server to process messages.

Each session corresponds to a separate listener thread and therefore controls the number of concurrently processed messages. Adjust this parameter when the JMS server does not fully use the available capacity of the machine and if you do not need to process messages in a specific message order.

Data type	Integer
Units	Sessions
Default	1
Range	1 through 2147483647
Recommended	<ul style="list-style-type: none"> • If you want to process messages in a strict message order, set the value to 1, so only one thread is ever processing messages. • If you want to process multiple messages simultaneously (known as “message concurrency”), set this property to a value greater than 1. Keep this value as low as possible to prevent overloading client applications. A good starting point for a 100% JMS workload with short transaction times is 2 to 4 sessions per processor. If longer running transactions exist, you may need more sessions, which should be determined by experimentation. • If you are using XA transactions, this property should always be set to 1. If it this property is set to a higher value, multiple messages are delivered in the same transaction, which is usually not the desired behavior.

Maximum retries:

The maximum number of times that the listener tries to deliver a message before the listener is stopped, in the range 0 through 2147483647.

The maximum number of times that the listener tries to deliver a message to a message-driven bean instance before the listener is stopped.

Data type	Integer
Units	Retry attempts
Default	0 (no retries)
Range	0 (no retries) through 2147483647

Maximum messages:

The maximum number of messages that the listener can process in one session with the JMS server.

For WebSphere embedded messaging or WebSphere MQ as the JMS provider, the listener processes all messages in the session as one batch within the same transaction. For a generic JMS provider, the listener processes each message in the session within a separate transaction.

Data type	Integer
Units	Number of messages
Default	1
Range	1 through 2147483647

Recommended

For WebSphere embedded messaging or WebSphere MQ as the JMS provider, if you want to process multiple messages in a single transaction, then set this value to more than 1. This enables multiple messages to be batch-processed into a single transaction, and eliminates much of the overhead of transactions on JMS messages.

CAUTION:

- If one message in the batch fails processing with an exception, the entire batch of messages is put back on the queue for processing.
- Any resource lock held by any of the interactions for the individual messages are held for the duration of the entire batch.
- Depending on the amount of processing that messages need, and if XA transactions are being used, setting a value greater than 1 can cause the transaction to time out. If an XA transaction does time out routinely because processing multiple messages exceeds the transaction timeout, reduce this property to 1 (to limit processing to one message per transaction) or increase your transaction timeout.

Configuring a listener port to handle late responses:

Use this task to configure the listener port extensions to handle late responses for an application that uses extended messaging.

To configure a listener port to handle late responses, use the administrative console to complete the following steps:

1. Define the properties of the listener port to be used for late responses, as described in Adding a new listener port or Configuring a listener port.
2. Configure the listener port extensions properties, to indicate that you want to handle late responses, as follows:
 - a. In the navigation pane, expand **Servers-> Application Servers-> application_server-> Extended Messaging Service-> Listener Port Extensions**
This displays a panel in the content pane for configuring the extensions properties of listener ports.
 - b. From the pull-down list, select the listener port that you want to configure to handle late responses.
 - c. Select the **Enable** checkbox to enable handling of late responses.
 - d. Optionally, change the values for the Request Interval and Request Timeout properties.
 - e. Click **OK**.

Adding a new listener port

Use this task to add a new listener port to the message listener service, so that message-driven beans can be associated with the port to retrieve messages.

To add a new listener port, use the administrative console to complete the following steps:

1. In the navigation pane, select **Servers-> Application Servers** This displays a table of the application servers in the administrative domain.
2. In the content pane, click the name of the application server. This displays the properties of the application server in the content pane.
3. In the Additional Properties table, select **Message Listener Service** This displays the Message Listener Service properties in the content pane.
4. In the content pane, select **Listener Ports**. This displays a list of the listener ports.

5. In the content pane, click **New**.
6. Specify appropriate properties for the listener port.
7. Click **OK**.
8. To save your configuration, click **Save** on the task bar of the Administrative console window.
9. To have the changed configuration take effect, stop then restart the application server.

If enabled, the listener port is started automatically when a message-driven bean associated with that port is installed.

Configuring a listener port

Use this task to change the properties of an existing listener port, used by message-driven beans associated with the port to retrieve messages.

To configure the properties of a listener port, use the administrative console to complete the following steps:

1. In the navigation pane, select **Servers-> Application Servers** This displays a table of the application servers in the administrative domain.
2. In the content pane, click the name of the application server. This displays the properties of the application server in the content pane.
3. In the Additional Properties table, select **Message Listener Service** This displays the Message Listener Service properties in the content pane.
4. In the content pane, click **Listener Ports**. This displays a list of the listener ports.
5. Click the listener port that you want to modify. This displays the properties of the listener port in the content pane.
6. Specify appropriate properties for the listener port.
7. Click **OK**.
8. To save your configuration, click **Save** on the task bar of the Administrative console window.
9. To have the changed configuration take effect, stop then restart the application server.

Deleting a listener port

Use this task to delete a listener port from the message listener service, to prevent message-driven beans associated with the port from retrieving messages.

To delete a listener port, use the administrative console to complete the following steps:

1. In the navigation pane, select **Servers-> Application Servers** This displays a table of the application servers in the administrative domain.
2. In the content pane, click the name of the application server. This displays the properties of the application server in the content pane.
3. In the Additional Properties table, select **Message Listener Service** This displays the Message Listener Service properties in the content pane.
4. In the content pane, select **Listener Ports**. This displays a list of the listener ports.
5. In the content pane, select the checkbox for the listener port that you want to delete.
6. Click **Delete**. This action stops the port (needed to allow the port to be deleted) then deletes the port.
7. To save your configuration, click **Save** on the task bar of the Administrative console window.
8. To have the changed configuration take effect, stop then restart the application server.

Configuring security for message-driven beans

Use this task to configure resource security and security permissions for message-driven beans.

Messages arriving at a listener port have no client credentials associated with them. The messages are anonymous.

To call secure enterprise beans from a message-driven bean, the message-driven bean needs to be configured with a RunAs Identity deployment descriptor. Security depends on the role specified by the RunAs Identity for the message-driven bean as an EJB component.

For more information about EJB security, see "Enterprise bean component security" in the Information Center. For more information about configuring security for your application, see "Assembling secured applications" in the Information Center.

JMS connections used by message-driven beans can benefit from the added security of using J2C container-managed authentication. To enable the use of J2C container authentication aliases and mapping, define a J2C container-managed alias on the JMS connection factory definition that the MDB is using to listen upon (defined by the **Connection factory JNDI name** property of the listener port). If defined, the listener uses the container-managed authentication alias for its JMSConnection security credentials instead of any application-managed alias. To set the container-managed alias, use the administrative console to complete the following steps:

1. To display the listener port settings, click **Servers-> application_server-> Message Listener Service-> Listener Ports-> listener_port**
2. To get the name of the JMS connection factory, look at the **Connection factory JNDI name** property.
3. Display the JMS connection factory properties. For example, to display the properties of a queue connection factory provided by the embedded WebSphere JMS provider, click **Resources-> WebSphere JMS Provider-> (In content pane, under Additional Properties) WebSphere Queue Connection Factories-> connection_factory**
4. Set the **Container-managed Authentication Alias** property.
5. Click **OK**

Administering listener ports

Use the following tasks to administer listener ports, which each define the association between a connection factory, a destination, and a message-driven bean.

You can use the WebSphere administrative console to administer listener ports, as described in the following tasks.

- Adding a new listener port
Use this task to create a new listener port, to specify a new association between a connection factory, a destination, and a message-driven bean. This enables deployed message-driven beans associated with the port to retrieve messages from the destination.
- Configuring a listener port
Use this task to view or change the configuration properties of a listener port.
- Starting a listener port
Use this task to start a listener port manually.
- Stopping a listener port
Use this task to stop a listener port manually.

Note: If configured as enabled, a listener port is started automatically when a message-driven bean associated with that port is installed. You do not normally need to start or stop a listener port manually.

Starting a listener port:

Use this task to start a listener port on an application server, to enable the listeners for message-driven beans associated with the port to retrieve messages.

A listener is active, that is able to receive messages from a destination, if the deployed message-driven bean, listener port, and message listener service are all started. Although you can start these components in any order, they must all be in a started state before the listener can retrieve messages.

If configured as enabled, a listener port is started automatically when a message-driven bean associated with that port is installed. However, you can start a listener port manually, as described in this topic.

When a listener port is started, the listener manager tries to start the listeners for each message-driven bean associated with the port. If a message-driven bean is stopped, the port is started but the listener is not started, and remains stopped. If you start a message-driven bean, the related listener is started.

To start a listener port on an application server, use the administrative console to complete the following steps:

1. If you want the listener for a deployed message-driven bean to be able to receive messages at the port, check that the message-driven bean has been started.
2. In the navigation pane, select **Servers-> Application Servers** This displays a table of the application servers in the administrative domain.
3. In the content pane, click the name of the application server. This displays the properties of the application server in the content pane.
4. In the Additional Properties table, select **Message Listener Service** This displays the Message Listener Service properties in the content pane.
5. In the content pane, select **Listener Ports**. This displays a list of the listener ports.
6. Select the checkbox for the listener port that you want to start.
7. Click **Start**.
8. To save your configuration, click **Save** on the task bar of the Administrative console window.

Stopping a listener port:

Use this task to stop a listener port on an application server, to prevent the listeners for message-driven beans associated with the port from retrieving messages.

When you stop a listener port as described in this topic, the listener manager stops the listeners for all message-driven beans associated with the port.

To stop a listener port on an application server, use the administrative console to complete the following steps:

1. In the navigation pane, select **Servers-> Application Servers** This displays a table of the application servers in the administrative domain.
2. In the content pane, click the name of the application server. This displays the properties of the application server in the content pane.
3. In the Additional Properties table, select **Message Listener Service** This displays the Message Listener Service properties in the content pane.
4. In the content pane, select **Listener Ports**. This displays a list of the listener ports.
5. In the content pane, select the listener port that you want to stop.
6. Click **Stop**.
7. To save your configuration, click **Save** on the task bar of the Administrative console window.
8. To have the changed configuration take effect, stop then restart the application server.

Important files for message-driven beans and extended messaging

The following files in the WAS_HOME/temp directory are important for the operation of the WebSphere Application Server messaging service, so should not be deleted. If you do need to delete the WAS_HOME/temp directory or other files in it, ensure that you preserve the following files.

server_name-durableSubscriptions.ser

You should not delete this file, because the messaging service uses it to keep track of durable subscriptions for message-driven beans. If you uninstall an application that contains a message-driven bean, this file is used to unsubscribe the durable subscription.

`server_name-AsyncMessageRequestLog.ser`

You should not delete this file, because the messaging service uses it to keep track of late responses that need to be delivered to the late response message handler for the extended messaging provider.

Troubleshooting message-driven beans

Use this overview task to help resolve a problem that you think is related to message-driven beans.

Message-driven beans support uses the standard WebSphere Application Server troubleshooting facilities. If you encounter a problem that you think might be related to the message-driven beans, complete the following stages:

1. Check for error messages about message-driven beans.

Check in the application server's SystemOut log at `was_home\logs\server\SystemOut` for error messages that indicate a problem with JMS resources for message-driven beans, such as listener ports.

The associated message reference information provides an explanation and any user actions to resolve the problem. (See "Message reference" in the Information Center.)

2. Check for more informational and error messages that might provide a clue to a related problem. If the JMS server is running, but you have problems accessing JMS resources, check for more error messages and extra details about the problem.

For messages related to WebSphere Messaging, look for the prefixes: MSGS and WMSG.

3. Check the Release Notes for specific problems and workarounds The section *Possible Problems and Suggested Fixes* of the Release Notes, available from the WebSphere Application Server library web site, is updated regularly to contain information about known defects and their workarounds. Check the latest version of the Release Notes for any information about your problem. If the Release Notes does not contain any information about your problem, you can also search the Technotes database on the WebSphere Application Server web site.
4. Check that message listener service has started. The message listener service is an extension to the JMS functions of the JMS provider. It provides a listener manager that controls and monitors one or more JMS listeners, which each monitor a JMS destination on behalf of a deployed message-driven bean.
5. Check your JMS resource configurations If the WebSphere Messaging functions seem to be running properly (the JMS server is running without problems), check that the JMS resources have been configured correctly. For example, check that the listener ports have been configured correctly and have been started.
6. Check for problems with the WebSphere Messaging functions For more information about troubleshooting WebSphere Messaging, see the related topics.
7. Get a detailed exception dump for messaging. If the information obtained in the preceding steps is still inconclusive, you can enable the application server debug trace for the "Messaging" group to provide a detailed exception dump.

Message-driven beans samples

The following examples are provided, as part of the WebSphere Samples Gallery, to illustrate use of the message-driven beans support. When the Samples are installed on your local machine, they are available to try out. Locate them at <http://localhost:9080/WSSamples/>. (The default port is 9080.) For more information about where to find the Samples Gallery, see "Samples Gallery" in the Information Center.

- Point-to-point samples:

- "Tutorial: Creating JMS message sample"

This tutorial is designed to help you develop and deploy a JMS message sample application that tests the WebSphere Application Server message-driven beans support in a point-to-point scenario. This sample illustrates how to develop and deploy an application that comprises the following components:

- A Java/JMS program that writes a message to a queue.
- A message-driven bean that is invoked by a JMS listener when a message arrives on a defined queue.

For more information about this sample, see the samples article "Tutorial: Creating JMS message sample" that is installed with the Samples option.

- "Sample: Message Listener (point-to-point)"

This sample is designed to demonstrate the use and behavior of message-driven beans for a simple point-to-point scenario. This sample uses the JMS message sample deployed in the sample above.

For more information about this sample, see the samples article "Sample: Message Listener (Point-to-Point)" that is installed with the Samples option.

- Publish/subscribe samples

- "Tutorial: Creating JMS message publish/subscribe sample"

This tutorial is designed to help you develop and deploy a JMS message sample application that tests the WebSphere Application Server message-driven beans support in a publish/subscribe scenario. This sample illustrates how to develop and deploy an application that comprises the following components:

- A client program that starts the message sequence by publishing a message to a selected topic.
- A message-driven bean that is invoked by a JMS listener when the broker passes a message to the listener from a topic to which it has subscribed.

For more information about this sample, see the samples article "Tutorial: Creating JMS message publish/subscribe sample" that is installed with the Samples option.

- "Sample: Message Listener (publish/subscribe)"

This sample is designed to demonstrate the use and behavior of message-driven beans for a simple publish/subscribe scenario. This sample uses the JMS message sample deployed in the publish/subscribe sample above.

For more information about this sample, see the samples article "Sample: Message Listener (publish/subscribe)" that is installed with the Samples option.

Chapter 7. Using application clients

An application client module is a Java ARchive (JAR) file that contains a client for accessing a Java application. Complete the following steps for developing different types of application clients.

1. Decide on a type of application client.
2. Develop the application client code.
 - a. Develop ActiveX application client code.
 - b. Develop J2EE application client code.
 - c. Develop pluggable application client code.
 - d. Develop thin application client code.

View the Samples gallery for more information about application clients. (See "Samples Gallery" in the Information Center.) Before you run the `basicCalculator` sample, verify that the Java Messaging Service (JMS) server is started.

These samples do not include client applications that communicate with Enterprise server-side examples.

Application clients

In a traditional client server environment, the client requests a service and the server fulfills the request. Multiple clients use a single server. Clients can also access several different servers. This model persists for Java clients except that now these requests use a client run-time environment.

In this model, the client application requires a servlet to communicate with the enterprise bean, and the servlet must reside on the same machine as the WebSphere Application Server.

With WebSphere Application Server Version 5, application clients now consist of the following models:

- ActiveX application client
- Applet client
- J2EE application client
- Pluggable and thin application clients

The *ActiveX application client* model, uses the Java Native Interface (JNI) architecture to programmatically access the Java virtual machine (JVM) API. Therefore the JVM code exists in the same process space as the ActiveX application (Visual Basic, VBScript, or Active Server Pages (ASP) files) and remains attached to the process until that process terminates.

In the *Applet client* model, a Java applet embeds in a HyperText Markup Language (HTML) document residing on a remote client machine from the WebSphere Application Server. With this type of client, the user accesses an enterprise bean in the WebSphere Application Server through the Java applet in the HTML document.

The *J2EE application client* is a Java application program that accesses enterprise beans, Java Database Connectivity (JDBC) APIs, and Java Message Service message queues. The J2EE application client program runs on client machines. This program follows the same Java programming model as other Java programs; however, the J2EE application client depends on the application client run time to configure its execution environment, and uses the Java Naming and Directory Interface (JNDI) name space to access resources.

The *pluggable and thin application clients* provide a lightweight Java client programming model. These clients are useful in situations where a Java client application exists but the application needs enhancements to use enterprise beans, or where the client application requires a thinner, more lightweight environment than the one offered by the J2EE application client. The difference between the thin

application client and the pluggable application client is that the thin application client includes a Java virtual machine (JVM) API, and the pluggable application client requires the user to provide this code. The pluggable application client uses the Sun Java Development Kit, and the thin application client uses the IBM Developer Kit for the Java platform.

The J2EE application client programming model provides the benefits of the J2EE platform for the Java client application. Use the J2EE application client to seamlessly develop, assemble, deploy and launch a client application. The tooling provided with the WebSphere platform supports the seamless integration of these stages to help the developer create a client application from start to finish.

When you develop a client application using and adhering to the J2EE platform, you can put the client application code from one J2EE platform implementation to another. The client application package can require redeployment using each J2EE platform deployment tool, but the code that comprises the client application remains the same.

The application client run time supplies a container that provides access to system services for the client application code. The client application code must contain a main method. The application client run time invokes this main method after the environment initializes and runs until the Java virtual machine code terminates.

The J2EE platform supports the application client use of *nicknames* or *short names*, defined within the client application deployment descriptor. These deployment descriptors identify enterprise beans or local resources (JDBC, Java Message Service (JMS), JavaMail and URL APIs) for simplified resolution through JNDI. This simplified resolution to the enterprise bean reference and local resource reference also eliminates changes to the client application code, when the underlying object or resource either changes or moves to a different server. When these changes occur, the application client can require redeployment.

The application client also provides initialization of the run-time environment for the client application. The deployment descriptor defines this unique initialization for each client application. The application client run time also provides support for security authentication to enterprise beans and local resources.

The application client uses the Java Remote Method Invocation-Internet InterORB Protocol (RMI-IIOP). Using this protocol enables the client application to access enterprise bean references and to use Common Object Request Broker Architecture (CORBA) services provided by the J2EE platform implementation. Use of the RMI-IIOP protocol and the accessibility of CORBA services assist users in developing a client application that requires access to both enterprise bean references and CORBA object references.

When you combine the J2EE and CORBA environments or programming models in one client application, you must understand the differences between the two programming models to use and manage each appropriately.

View the Samples gallery for more information about application clients. (See "Samples Gallery" in the Information Center.) Before you run the `basicCalculator` sample, verify that the Java Messaging Service (JMS) server is started.

These samples do not include client applications that communicate with Enterprise server-side examples.

Application client functions

Use the following table to identify the available functions in the different types of clients.

Available functions	ActiveX client	Applet client	J2EE client	Pluggable client	Thin client
Provides all the benefits of a J2EE platform	Yes	No	Yes	No	No

Portable across all J2EE platforms	No	No	Yes	No	No
Provides the necessary run-time support for communication between a client and a server	Yes	Yes	Yes	Yes	Yes
Supports the use of nicknames in the deployment descriptor files. Note: Although you can edit deployment descriptor files, do not use the administrative console to modify them.	Yes	No	Yes	No	No
Supports use of the RMI-IIOP protocol	Yes	Yes	Yes	Yes	Yes
Browser-based application	No	Yes	No	No	No
Enables development of client applications that can access enterprise bean references and CORBA object references	Yes	Yes	Yes	Yes	Yes
Enables the initialization of the client application run-time environment	Yes	No	Yes	No	No
Supports security authentication to enterprise beans (see "Configuring authentication mechanisms" in the Information Center)	Yes	Limited	Yes	Yes	Yes
Supports security authentication to local resources (see "Configuring authentication mechanisms" in the Information Center)	Yes	No	Yes	No	No
Requires distribution of application to client machines	Yes	No	Yes	Yes	Yes
Enables access to enterprise beans and other Java classes through Visual Basic, VBScript, and Active Server Pages (ASP) code	Yes	No	No	No	No
Provides a lightweight client suitable for download	No	Yes	No	Yes	Yes
Enables access JNDI APIs for enterprise bean resolution	Yes	Yes	Yes	Yes	Yes
Runs on client machines that use the Sun Java Runtime Environment	No	No	No	Yes	No
Supports CORBA services (using CORBA services can render the application client code nonportable)	No	No	Yes	No	No

ActiveX application clients

WebSphere Application Server provides an ActiveX to EJB bridge that enables ActiveX programs to access enterprise beans through a set of ActiveX automation objects.

The bridge accomplishes this access by loading the Java virtual machine (JVM) into any ActiveX automation container such as Visual Basic, VBScript, and Active Server Pages (ASP).

There are two main environments in which the ActiveX to EJB bridge runs:

- **Client applications**, such as Visual Basic and VBScript, are programs that a user starts from the command line, desktop icon, or Start menu shortcut.
- **Client services**, such as Active Server Pages, are programs started by some automated means like the Services control panel applet.

The ActiveX to EJB bridge uses the Java Native Interface (JNI) architecture to programmatically access the JVM code. Therefore the JVM code exists in the same process space as the ActiveX application (Visual Basic, VBScript, or ASP) and remains attached to the process until that process terminates. To create JVM code, an ActiveX client program calls the XJBInit() method of the XJB.JClassFactory object. For more information about creating JVM code for an ActiveX program, see ActiveX to EJB bridge, initializing JVM code.

After an ActiveX client program has initialized the JVM code, the program calls several methods to create a proxy object for the Java class. When accessing a Java class or object, the real Java object exists in the JVM code; the automation container contains the proxy for that Java object. The ActiveX program can use the proxy object to access the Java class, object fields, and methods. For more information about using Java proxy objects, see ActiveX to EJB bridge, using Java proxy objects. For more information about calling methods and access fields, see ActiveX to EJB bridge, calling Java methods and ActiveX to EJB bridge, accessing Java fields.

The client program performs primitive data type conversion through the COM IDispatch interface (use of the IUnknown interface is not directly supported). Primitive data types are automatically converted between native automation types and Java types. All other types are handled automatically by the proxy objects. For more information about data type conversion, see ActiveX to EJB bridge, converting data types.

Any exceptions thrown in Java code are encapsulated and thrown again as a COM error, from which the ActiveX program can determine the actual Java exceptions. For more information about handling exceptions, see ActiveX to EJB bridge, handling errors.

The ActiveX to EJB bridge supports both free-threaded and apartment-threaded access and implements the free threaded marshaler (FTM) to work in a hybrid environment such as Active Server Pages. For more information about the support for threading, see ActiveX to EJB bridge, using threading.

Applet clients

The applet client provides a browser-based Java run time capable of interacting with enterprise beans directly, instead of indirectly through a servlet.

This client is designed to support users who want a browser-based Java client application programming environment that provides a richer and more robust environment than the one offered by the **Applet > Servlet > enterprise bean** model.

The programming model for this client is a hybrid of the Java application thin client and a servlet client. When accessing enterprise beans from this client, the applet can consider the enterprise bean object references as CORBA object references.

No tooling support exists for this client to develop, assemble or deploy the applet. You are responsible for developing the applet, generating the necessary client bindings for the enterprise beans and CORBA

objects, and bundling these pieces together to install or download to the client machine. The Java applet client provides the necessary run time to support communication between the client and the server. The applet client run time is provided through the Java applet browser plug-in that you install on the client machine using the WebSphere Application Server Client CD.

Generate client-side bindings using the Application Assembly Tool **5.1+** Chapter 29, “Assembling applications with the Assembly Toolkit,” on page 1391. An applet can utilize these bindings, or you can generate client-side bindings using the **rmic** command. This command is part of the IBM Developer Kit, Java edition that is installed with the WebSphere Application Server.

The applet client uses the RMI-IIOP protocol. Using this protocol enables the applet to access enterprise bean references and CORBA object references, but the applet is restricted in using some supported CORBA services.

If you combine the enterprise bean and CORBA environments in one applet, you must understand the differences between the two programming models, and you must use and manage each model appropriately.

The applet environment restricts access to external resources from the browser run-time environment. You can make some of these resources available to the applet by setting the correct security policy settings in the WebSphere Application Server `client.policy` file. If given the correct set of permissions, the applet client must explicitly create the connection to the resource using the appropriate API. This client does not perform initialization of any service that the client applet can need. For example, the client application is responsible for the initialization of the naming service, either through the CosNaming, or the Java Naming and Directory Interface (JNDI) APIs.

J2EE application clients

The J2EE application client programming model provides the benefits of the Java 2 Platform for WebSphere Application Server Enterprise product.

The J2EE platform offers the ability to seamlessly develop, assemble, deploy and launch a client application. The tooling provided with the WebSphere platform supports the seamless integration of these stages to help the developer create a client application from start to finish.

When you develop a client application using and adhering to the J2EE platform, you can put the client application code from one J2EE platform implementation to another. The client application package can require redeployment using each J2EE platform deployment tool, but the code that comprises the client application does not change.

The J2EE application client run time supplies a container that provides access to system services for the application client code. The J2EE application client code must contain a main method. The J2EE application client run time invokes this main method after the environment initializes and runs until the Java virtual machine application terminates.

Application clients can use *nicknames* or *short names*, defined within the client application deployment descriptor with the J2EE platform. These deployment descriptors identify enterprise beans or local resources (Java Database Connectivity (JDBC), Java Message Service (JMS), JavaMail and URL APIs) for simplified resolution through JNDI use. This simplified resolution to the enterprise bean reference and local resource reference also eliminates changes to the application client code, when the underlying object or resource either changes or moves to a different server. When these changes occur, the application client can require redeployment. Although you can edit deployment descriptor files, do not use the administrative console to modify them.

The J2EE application client also provides initialization of the run-time environment for the client application. The deployment descriptor defines this unique initialization for each client application. The J2EE application client run time also provides support for security authentication to the enterprise beans and local resources.

The J2EE application client uses the Java Remote Method Invocation technology run over Internet Inter-Orb Protocol (RMI-IIOP). Using this protocol enables the client application to access enterprise bean references and to use Common Object Request Broker Architecture (CORBA) services provided by the J2EE platform implementation. Use of the RMI-IIOP protocol and the accessibility of CORBA services assist users in developing a client application that requires access to both enterprise bean references and CORBA object references.

When you combine the J2EE and the CORBA WebSphere Application Server Enterprise environments or programming models in one client application, you must understand the differences between the two programming models to use and manage each appropriately.

Pluggable application clients

The pluggable application client provides a lightweight, downloadable Java application run time capable of interacting with enterprise beans.

The pluggable application client requires that you have previously installed the Sun Java Runtime Environment (JRE) files. In all other aspects, the pluggable application client, and the thin application client are similar.

Note: The pluggable client is only available on the Windows platform.

This client is designed to support those users who want a lightweight Java client application programming environment, without the overhead of the J2EE platform on the client machine. The programming model for this client is heavily influenced by the CORBA programming model, but supports access to enterprise beans.

When accessing enterprise beans from this client, the client application can consider the enterprise beans object references as CORBA object references.

Tooling does not exist on the client; however, tooling does exist on the server. You are responsible for developing the client application, generating the necessary client bindings for the enterprise bean and CORBA objects, and after bundling these pieces together, installing them on the client machine.

The pluggable application client provides the necessary run time to support the communication needs between the client and the server.

The pluggable application client uses the RMI-IIOP protocol. Using this protocol enables the client application to access enterprise bean references and CORBA object references and use any supported CORBA services. Using the RMI-IIOP protocol along with the accessibility of CORBA services can assist a user in developing a client application that needs to access both enterprise bean references and CORBA object references.

When you combine the J2EE and CORBA environments in one client application, you must understand the differences between the two programming models to use and manage each appropriately.

The pluggable application client run time provides the necessary support for the client application for object resolution, security, Reliability Availability and Serviceability (RAS), and other services. However, this client does not support a container that provides easy access to these services. For example, no support exists for using *nicknames* for enterprise beans or local resource resolution. When resolving to an

enterprise bean (using either the Java Naming and Directory Interface (JNDI) API or CosNaming) sources, the client application must know the location of the name server and the fully qualified name used when the reference was bound into the name space.

When resolving to a local resource, the client application cannot resolve to the resource through a JNDI lookup. Instead the client application must explicitly create the connection to the resource using the appropriate API (JDBC, Java Message Service (JMS), and so on). This client does not perform initialization of any of the services that the client application might require. For example, the client application is responsible for the initialization of the naming service, either through CosNaming or JNDI APIs.

The pluggable application client offers access to most of the available client services in the J2EE application client. However, you cannot access the services in the pluggable client as easily as you can in the J2EE application client. The J2EE client has the advantage of performing a simple Java Naming and Directory Interface (JNDI) name space lookup to access the desired service or resource. The pluggable client must code explicitly for each resource in the client application. For example, looking up an enterprise bean Home object requires the following code in a J2EE application client:

```
        java.lang.Object ejbHome = initialContext.lookup("java:/comp/env/ejb/MyEJBHome"
);
    MyEJBHome = (MyEJBHome)javax.rmi.PortableRemoteObject.narrow(ejbHome,
MyEJBHome.class);
```

However, you need more explicit code in a Java pluggable application client:

```
        java.lang.Object ejbHome = initialContext.lookup("the/fully/qualified
/path/to/actual/home/in/namespace/MyEJBHome");
    MyEJBHome = (MyEJBHome)javax.rmi.PortableRemoteObject.narrow(ejbHome,
MyEJBHome.class);
```

In this example, the J2EE application client accesses a logical name from the `java:/comp` name space. The J2EE client run time resolves that name to the physical location and returns the reference to the client application. The pluggable client must know the fully qualified physical location of the enterprise bean Home object in the name space. If this location changes, the pluggable client application must also change the value placed on the `lookup()` statement.

In the J2EE client, the client application is protected from these changes because it uses the logical name. A change can require a redeployment of the EAR file, but the actual client application code remains the same.

The pluggable application client is a traditional Java application that contains a *main* function. The WebSphere pluggable application client provides run-time support for accessing remote enterprise beans, and provides the implementation for various services (security, Workload Management (WLM), and others). This client can also access CORBA objects and CORBA-based services. When using both environments in one client application, you need to understand the differences between the enterprise bean and the CORBA programming models to manage both environments.

For instance, the CORBA programming model requires the CORBA CosNaming name service for object resolution in a name space. The enterprise beans programming model requires the JNDI name service. The client application must initialize and properly manage these two naming services.

Another difference applies to the enterprise bean model. Use the Java Naming and Directory Interface (JNDI) implementation in the enterprise bean model to initialize the Object Request Broker (ORB). The client application is unaware that an ORB is present. The CORBA model, however, requires the client application to explicitly initialize the ORB through the `ORB.init()` static method.

The pluggable application client provides a batch command that you can use to set the `CLASSPATH` and `JAVA_HOME` environment variables to enable the pluggable application client run time.

Thin application clients

The thin application client provides a lightweight, downloadable Java application run time capable of interacting with enterprise beans.

This client is designed to support those users who want a lightweight Java client application programming environment, without the overhead of the J2EE platform on the client machine. The programming model for this client is heavily influenced by the CORBA programming model, but supports access to enterprise beans.

When accessing enterprise beans from this client, the client application can consider the enterprise beans object references as CORBA object references.

Tooling does not exist on the client, it exists on the server. You are responsible for developing the client application, generating the necessary client bindings for the enterprise bean and CORBA objects, and bundling these pieces together to install on the client machine.

The thin application client provides the necessary run-time to support the communication needs between the client and the server.

The thin application client uses the RMI-IIOP protocol. Using this protocol enables the client application to access not only enterprise bean references and CORBA object references, but also allows the client application to use any supported CORBA services. Using the RMI-IIOP protocol along with the accessibility of CORBA services can assist a user in developing a client application that needs to access both enterprise bean references and CORBA object references.

When you combine the J2EE and CORBA environments in one client application, you must understand the differences between the two programming models, to use and manage each appropriately.

The thin application client run time provides the necessary support for the client application for object resolution, security, Reliability Availability and Servicability (RAS), and other services. However, this client does not support a container that provides easy access to these services. For example, no support exists for using *nicknames* for enterprise beans or local resource resolution. When resolving to an enterprise bean (using either Java Naming and Directory Interface (JNDI) or CosNaming) sources, the client application must know the location of the name server and the fully qualified name used when the reference was bound into the name space. When resolving to a local resource, the client application cannot resolve to the resource through a JNDI lookup. Instead the client application must explicitly create the connection to the resource using the appropriate API (JDBC, Java Message Service (JMS), and so on). This client does not perform initialization of any of the services that the client application might require. For example, the client application is responsible for the initialization of the naming service, either through CosNaming or JNDI APIs.

The thin application client offers access to most of the available client services in the J2EE application client. However, you cannot access the services in the thin client as easily as you can in the J2EE application client. The J2EE client has the advantage of performing a simple Java Naming and Directory Interface (JNDI) name space lookup to access the desired service or resource. The thin client must code explicitly for each resource in the client application. For example, looking up an enterprise bean Home requires the following code in a J2EE application client:

```
java.lang.Object ejbHome = initialContext.lookup("java:/comp/env/ejb/MyEJBHome");
MyEJBHome = (MyEJBHome)javax.rmi.PortableRemoteObject.narrow(ejbHome, MyEJBHome.class);
```

However, you need more explicit code in a Java thin application client:

```
java.lang.Object ejbHome =
initialContext.lookup("the/fully/qualified/path/to/actual/home/in/namespac/MyEJBHome");
MyEJBHome = (MyEJBHome)javax.rmi.PortableRemoteObject.narrow(ejbHome, MyEJBHome.class);
```

In this example, the J2EE application client accesses a logical name from the `java:/comp` name space. The J2EE client run time resolves that name to the physical location and returns the reference to the client application. The thin client must know the fully qualified physical location of the enterprise bean Home in the name space. If this location changes, the thin client application must also change the value placed on the `lookup()` statement.

In the J2EE client, the client application is protected from these changes because it uses the logical name. A change might require a redeployment of the EAR file, but the actual client application code remains the same.

The thin application client is a traditional Java application that contains a *main* function. The WebSphere thin application client provides run-time support for accessing remote enterprise beans, and provides the implementation for various services (security, Workload Management (WLM), and others). This client can also access CORBA objects and CORBA based services. When using both environments in one client application, you need to understand the differences between the enterprise bean and CORBA programming models to manage both environments.

For instance, the CORBA programming model requires the CORBA CosNaming name service for object resolution in a name space. The enterprise beans programming model requires the JNDI name service. The client application must initialize and properly manage these two naming services.

Another difference applies to the enterprise bean model. Use the Java Naming and Directory Interface (JNDI) implementation in the enterprise bean model to initialize the Object Request Broker (ORB). The client application is unaware that an ORB is present. The CORBA model, however, requires the client application to explicitly initialize the ORB through the `ORB.init()` static method.

The thin application client provides a batch command that you can use to set the `CLASSPATH` and `JAVA_HOME` environment variables to enable the thin application client run time.

Migration tips for application clients

Tips for migrating thin application client code:

The Java invocation used to run non-J2EE application clients has changed in Version 5.0. You must specify `-Xbootclasspath/p:%WAS_BOOTCLASSPATH%` on Windows systems or `-Xbootclasspath/p:$WAS_BOOTCLASSPATH` on UNIX systems when you invoke the Java command. Set the `WAS_BOOTCLASSPATH` environment variable in one of the following:

- `setupClient.bat` for Windows systems or `setupClient.sh` for UNIX systems on a WebSphere Application Server client installation.
- `setupCmdLine.bat` for Windows systems or `setupCmdLine.sh` for UNIX systems on a WebSphere Application Server installation.

For more information about using `-Xbootclasspath`, view sample code at the following path after you preform the application client installation:

```
install_root\samples\bin\ActiveXBridgeClients\VB\XJBExamples\modXJBHelpers.bat
```

Tips for migrating J2EE application client code:

If your J2EE application client uses resource references and you configured those resources using the Application Client Resource Configuration Tool (ACRCT), you must run the **ClientUpgrade** command to migrate the resource configuration information in WebSphere Application Server Version 5.

Installing application clients

All client applications run on a machine with the WebSphere Application Server installed. However, if the system does not have the Application Server installed, you can install Application Server clients, which provide a stand-alone client run-time environment for your client applications. See the Supported Prerequisites page on the IBM external Web site for more information on supported product platforms.

This article describes how to install the WebSphere Application Server clients using the installation image on the product CD-ROM labelled, **Application Clients**. The steps that follow provide enough detail to guide you through preparing for, choosing, and installing the variety of options and features provided. To prepare for installation and to make yourself familiar with installation options, complete the steps in this article and read the related topics, before you start to use the installation tools. Specifically, read these topics before installing the product:

- Tips for installing the embedded messaging feature
- Installing silently
- Best practices for installing

As a general rule, if you launch an installation and there is a problem such as not having enough temporary space or not having the right packages on your Linux or UNIX-based systems, then cancel the installation, make the required changes, and restart the installation to pick up changes you made.

Although it is not supported or recommended, you can install this product as a non-root user on a UNIX operating system, or from a user ID that is not part of the Administrator group on a Windows platform. However, there are certain components, such as the Embedded Messaging Client feature, that require you to install as root or as part of the Administrator group.

As previously mentioned, this installation method is not supported or recommended, but you can install the application client product on a machine with WebSphere Application Server installed. However, there are certain components, such as the Embedded Messaging Client feature, which might not install if the feature has already been installed during the WebSphere Application Server installation. On Windows platforms, WebSphere Application Server clients assume that the Embedded Messaging Client is installed in its default location; whereas, the WebSphere Application Server installation can install the messaging client in a different location.

1. Prepare a Linux or UNIX operating platform for the Embedded Messaging Client feature.

If you are installing the embedded messaging feature, you must create two operating system groups as described in Installing WebSphere embedded messaging as the JMS provider.

The Solaris Operating Environment and HP-UX also require you to increase kernel settings as described in Installing WebSphere embedded messaging as the JMS provider.

For other platform-specific information about using the embedded messaging feature, see "Tips for installing the embedded messaging feature" in the Information Center.

2. Start the installation.

- a. Issue `launchpad.sh` (or `launchpad.bat`) to initiate the launchpad tool and begin the installation process.

- b. Click **Install the product** from the launchpad tool to launch the InstallShield for MultiPlatforms installation wizard. This action launches the installation wizard.

The Readme documentation to which the launchpad links is the `readme.html` file in the CD root directory. The `readme` directory off the root of the CD has more detailed Readme files. The Installation Guide is in the `/docs` directory of the CD root directory.

When you install application clients, the current working directory must be the directory where the installer binary program is located. This placement is important because the resolution of the class files location is done in the current working directory. For example:

```
cd /install_root
./install
```

or when installing from the product CD-ROM:

```
cd <CD mount point>
./install
```

Failing to use the correct working directory can cause ISMP errors that abort the installation.

The installation wizard does not upgrade or remove previous WebSphere Application Server clients installation automatically. However, you must uninstall any previous installation manually or the installation wizard aborts the installation.

- c. As indicated in the previous example, you can start the installation wizard from the product CD-ROM, using the command line. The installation program is in the operating system platform directory on the product CD-ROM.

On other Linux platforms and UNIX-based platforms, run the `./install` command.

On Windows platforms, run the `Install.exe` command.

- d. You can also perform a silent installation using the `-options responsefile` parameter, which causes the installation wizard to read your responses from the options response file, instead of from the interactive graphical user interface. Customize the response file before installing silently. After customizing the file, issue the command to silently install. Silent installation is particularly useful if you install the product often.

The rest of this procedure assumes that you are using the installation wizard. There are corresponding entries in the response file for every prompt that is described as part of the wizard. Review the description of the response file for more information. Comments in the file describe how to customize its options.

3. Click **Next** to continue when the Welcome panel is displayed. The license displayed during the GUI installation can contain characters that display incorrectly in Japanese. For example, the section labeled Part 1 does not show the number 1. These missing characters do not significantly affect the content of the license agreement.
 - a. Click the radio button beside the **I accept the terms in the license agreement** message if you agree to the license agreement, and click **Next** to continue.

4. Choose a type of installation, and click **Next**.

If you use the GUI, you can choose a Typical installation type, which installs J2EE and Java Thin Application Client, Samples and Embedded Messaging Client features, or a Custom installation type.

The Custom installation type lets you select which features to install. However you can not install the J2EE and Java Thin Application Client feature and the Pluggable Application Client feature together.

(Windows Only) If you select the **ActiveX to EJB Bridge** feature, then the following is displayed in a dialog box: Do you want to add Java runtime to the system path and make it the default JRE? If you answer **Yes**, then the Java run time is added to the beginning of the system path. If you answer **No**, then the ActiveX to EJB Bridge does not function from the Active Server Pages (ASP), unless you add the Java run time to the path. To add the Java run time later, see the topic **ActiveX application clients** or **reinstall Application clients**.(Windows Only) If you select the **Applet Client** feature, then the following message might be displayed: An existing JDK or JRE has been detected on your computer. You chose to install the Applet Client, which will overwrite the registry entries for this JDK or JRE. Do you want to continue and install the Applet Client? If you select **Yes**, the installation overrides the registry on your machine. If you select **No**, the applet client feature is not installed, and you are directed back feature dialog box.

5. Install the samples development environment. If you choose to install any of the Samples features, a message box is displayed requesting conformation to install the Samples development environment.
 - a. Click **Yes** to install the Samples development environment that includes the IBM Developer Kit and the Apache Ant Tool.
 - b. Click **No** to skip installing the Samples development environment.
6. Specify a destination directory. Click **Next** to continue.

- a. Ensure that there is adequate space available in the target directory.
 - b. Specify a target directory for the WebSphere Application Server clients product. If you install the Embedded Messaging Client feature, then you cannot change its default installation directory.
 - c. Enter the required target directory to proceed to the next panel. Deleting the default target location and leaving an installation directory field empty prevents you from continuing the installation process.
7. Enter the host name of the WebSphere Application Server machine. Click **Next** to continue. If you are connecting to a product Version 4 server or you are not using the default port, you must specify the server port number. The default port number for product Version 4 server is 900, and the default port number for product Version 5 is 2809.
 8. Review the summary information, and click **Next** to install the product code or you might also click **Back** to change your specifications. When the installation is complete, the wizard displays the `install_root\logs\mq_install.log` installation log, if you selected the Embedded Messaging Client feature, and there are errors with its installation.
 9. Review the `mq_install.log` installation log, if it appears. Click **Next** to continue.
 10. Click **Finish** to exit the wizard, after the WebSphere Application Server client installs.

You successfully installed WebSphere Application Server clients and the features you selected.

If the installation is not successful, fix the error as indicated in the installation error message. For example, if you do not have enough disk space, add more space, and reinstall application clients.

Best practices for installing application clients

The following table offers tips for installing application clients on multiple platforms.

Operating environment	Tip
Linux and UNIX systems	Spaces are not supported in the name of the installation directory on Linux and UNIX platforms.
UNIX systems	When client application installations are successful, the return code 1 is issued from the UNIX shell where you issued the <code>/install</code> command. Any other return code indicates an unsuccessful installation.
Solaris systems	Double-byte character set (DBCS) characters are not supported in the name of the installation directory on Solaris systems.
Windows NT systems	Spaces are not supported in the name of the installation directory. Note: WebSphere Application Server Version 5.1 does not support the Windows NT platform.
All platforms	Reserve at least 4 to 5MB free space in the target platform temporary directory.
All platforms	When specifying a different temporary directory while installing application clients, the following message is displayed if the target platform default temporary directory does not have enough free space to install application clients: <pre>Error writing file = There may not be enough temporary disk space. Try using -is:tempdir to use a temporary directory on a partition with more disk space.</pre> <p>Use the <code>-is:tempdir</code> installation option to specify a different temporary directory. For example, the following command uses <code>/swap</code> as a temporary directory during installation:</p> <pre>./install -is:tempdir /swap</pre>

All platforms	<p>After the installation, when changing the installation settings for the WebSphere Application Server host name and the port number, edit the <code>setupClient.bat</code> for Windows or <code>setupClient.sh</code> for UNIX. Change the <code>DEFAULTSERVERNAME</code> and <code>SERVERPORTNUMBER</code> to the new WebSphere Application Server host name and port number, respectively. If the <code>SERVERPORTNUMBER</code> is not set, then the default is 2809. Review the following example:</p> <pre>set DEFAULTSERVERNAME=NDServerName set SERVERPORTNUMBER=9810</pre> <p>The <code>setupClient.bat</code> file or <code>setupClient.sh</code> file is located in the <code>bin</code> sub-directory under the application clients installation destination.</p>
---------------	---

Installing application clients silently

Use these steps to perform a silent installation, which uses the installation wizard to install the product. Instead of displaying a user interface, the silent installation provides interaction between you and the wizard by reading all of your responses from a file that you must customize.

1. Ensure that the user ID that you are using to run the silent installation has sufficient authority to perform the task.

If you are installing the embedded messaging feature, you must create two operating system groups as described in the article, [Installing WebSphere embedded messaging as the JMS provider](#).

Although it is not supported or recommended, you can install this product as a non-root user on a UNIX-based operating system, or from a user ID that is not apart of the Administrator group on a Windows platform. However, there are certain components, such as the Embedded Messaging Client feature, that require you to install as root or as part of the Administrator group.

2. Customize the option response file.
 - a. Locate the sample options response file. The file name is `setup.response` in the operating system platform directory on the product CD-ROM.
 - b. Make a copy to preserve the original response file. For example, copy the file as `myoptionsfile`.
 - c. Edit the copy in your flat file editor of choice, on the target operating system. Read the directions within the response file to choose appropriate values.

Note: To prepare the file for a silent installation on AIX, use UNIX line-end characters (0x0D0A) to terminate each line of the options response file.

- d. Make the first non-commented option `-silent` to have a silent install.
 - e. Include custom option responses that reflect parameters for your system.
 - f. Follow the instructions in the response file to choose appropriate values.
 - g. Save the file.
3. Issue a command to use your custom response file: `Install.exe -options myoptionsfile` for Windows platforms `install -options ./myoptionsfile` for Linux and UNIX platforms
The sample options response file is located in the operating-system platform directory on the product CD-ROM.
 4. Optional: Restart your machine in response to the prompt that appears on Windows platforms. If you install the Embedded Messaging Client feature, `-P JMSsupport.active="true"`, certain conditions, such as a locked file, might require you to restart your system. You have the option of restarting immediately, after which the installation program resumes the installation process. You can also defer restarting to a more convenient time, such as after the installation is complete.

You installed application clients silently by using the response file.

To further verify that the silent installation was successful, examine the `WAS.Client.Install.log` file for a line similar to:

```
The InstallShield Wizard has successfully installed IBM WebSphere Application Server clients, Version 5.1.
```

If you installed the Embedded Messaging Client feature, you can search the `mq_install.log` for any errors.

Developing application clients

This topic provides the steps for programming application clients to access resource objects defined on the server.

To use application clients to access a remote object on the server, develop your application clients as described in the following steps:

1. Create an instance of the object that you want to access from the remote server.
2. Specify the user ID and password on the connection method, when you create a connection to the server. Security must be enabled.
3. Assemble the application client `.ear` file using the Assembly Toolkit. Assemble the application client `.ear` file on any development machine where the Assembly Toolkit is installed.
4. Add the resource to the client deployment descriptor by completing the binding JNDI name for the resource object on the server.
5. Distribute the configured `.ear` file to the client machines.
6. Deploy the application client.
7. Configure the application client resources.

After you develop the application client code, run the application client.

Developing ActiveX application client code

This topic provides an outline for developing an ActiveX program, such as Visual Basic, VBScript, and Active Server Pages, to use the WebSphere ActiveX to EJB bridge to access enterprise beans.

This topic assumes that you are familiar with ActiveX programming. Consider the information given in ActiveX to EJB bridge as good programming guidelines.

To use the ActiveX to EJB bridge to access a Java class, develop your ActiveX program to complete the following steps:

1. Create an instance of the `XJB.JClassFactory` object.
2. Create Java Virtual Machine (JVM) code within the ActiveX program process, by calling the `XJBInit()` method of the `XJB.JClassFactory` object. After the ActiveX program has created an `XJB.JClassFactory` object and called the `XJBInit()` method, the JVM code is initialized and ready for use.
3. Create a proxy object for the Java class, by using the `XJB.JClassFactory FindClass()` and `NewInstance()` methods. The ActiveX program can use the proxy object to access the Java class, object fields, and methods.
4. Call methods on the Java class, using the Java method invocation syntax, and access Java fields as required.
5. Use the helper functions to do the conversion in cases where automatic conversion is not possible. You can convert between the following data types:
 - Java Byte and Visual Basic Byte
 - Visual Basic Currency types and Java 64-bit

6. Implement methods to handle any errors returned from the Java class. In Visual Basic or VBScript, use the Err.Number and Err.Description fields to determine the actual Java error.

After you develop the ActiveX client code, start the ActiveX application.

Starting an ActiveX application

To run an ActiveX client application that is to use the ActiveX to Enterprise Java Beans (EJB) bridge, you must perform some initial configuration to set appropriate environment variables and to enable the ActiveX to EJB bridge to find its XJB.JAR file and the Java run time. This initial configuration sets up the environment within which the ActiveX client application can run.

To perform the required configuration, complete one or more of the following subtasks:

1. Start an ActiveX application and configure service programs.
2. Start an ActiveX application and configuring non-service programs

Starting an ActiveX application and configuring service programs

To run an ActiveX service program such as Active Server Page (ASP) that is to use the ActiveX to the Enterprise Java Bean (EJB) bridge, you must perform some initial configuration to set appropriate environment variables and to enable the ActiveX to EJB bridge to find its XJB.JAR file and the Java run time. This configuration sets up the environment within which the ActiveX service program can run.

The XJB.JClassFactory must find the Java run time dynamic link library (DLL) when initializing. In a service program such as Internet Information Server you cannot specify a path for its processes independently; you must set the process paths in the system PATH variable. This limitation means that you can only have a single Java Virtual Machine (JVM) version available on a machine using ASP.

To add the Java Runtime Environment (JRE) directories to your system path, complete one of the following subtasks:

1. On Windows 2000 systems, complete the following substeps:
 - a. Open the Control Panel, then double-click the **System** icon.
 - b. Click the **Advanced** tab on the System Properties window.
 - c. Click **Environment Variables**.
 - d. Edit the Path variable in the System Variables window.
 - e. Add the following information to the beginning of the path that is displayed in the Variable Value field:

`C:\WebSphere\AppClient\Java\jre\bin;C:\WebSphere\AppClient\Java\jre\bin\classic;`

where C:\WebSphere\AppClient is the directory in which you installed the Java client in the WebSphere product.

- f. Click **OK** in the Edit System Variable window to apply the changes.
- g. Click **OK** in the Environment Variables window.
- h. Click **OK** in the System Properties window.
- i. Restart Windows 2000.

2. On Windows NT, complete the following substeps:

Note: WebSphere Application Server Version 5.1 does not support the Windows NT platform.

- a. Open the Control Panel, then double-click the System icon.
- b. Click the Environment tab on the System Properties window.
- c. In the System Variables window, edit the Path variable.
- d. Add the following to the beginning of the path displayed in the Value input box:

`C:\WebSphere\AppClient\Java\jre\bin;C:\WebSphere\AppClient\Java\jre\bin\classic;`

Where C:\WebSphere\AppClient is the directory in which you installed the Java client in the WebSphere product.

- e. Click **Set** to apply the changes.
- f. Click **OK**.
- g. Restart your Windows NT system.

After you change the system PATH variable you must reboot the Internet Information Server machine so that Internet Information Server can see the change.

Starting an ActiveX application and configuring non-service programs

To run an ActiveX program initiated from an icon or command line (a non-service program) that is to use the ActiveX to the Enterprise Java Beans (EJB) bridge, you must perform some initial configuration to set appropriate environment variables and to enable the ActiveX to EJB bridge to find its XJB.JAR file and the Java run-time environment. This uses a batch file to set up the environment within which the ActiveX program can run.

To perform the required configuration, complete the following steps:

1. Edit the setupCmdLineXJB.bat file to specify appropriate values for the environment variables required by the ActiveX to EJB bridge. For more information about these environment variables, see ActiveX to EJB bridge, environment and configuration. For more information about creating a JVM for an ActiveX program, see ActiveX to EJB bridge, initializing the Java Virtual Machine (JVM). After the ActiveX program has created an XJB.JClassFactory object and called the XJBInit() method, the JVM is initialized and ready for use.
2. Start the ActiveX client application by using one of the following methods:
 - Use the launchClientXJB.bat file to start the application. For example:


```
launchClientXJB MyApplication.exe parm1 parm2
```

or

```
launchClientXJB MyApplication.vbp
```
 - Use the setupCmdLineXJB.bat file to create an environment in which to run the application, then start the application from within that environment.

setupCmdLineXJB.bat, launchClientXJB.bat and other ActiveX batch files

This topic provides reference information about the aids that client applications and client services can use to access the ActiveX to EJB bridge. These enable the ActiveX to Enterprise Java Beans (EJB) bridge to find its XJB.JAR file and the Java run-time environment.

Location

The include file is located in the *was_client_home*\aspIncludes directory. You can include the file into your Active Server Pages (ASP) application with the following syntax in your ASP page:

```
<-- #include virtual ="/WSASPIncludes/setupASPXJB.inc" -->
```

This syntax assumes that you have created a virtual directory in Internet Information Server called WSASPIncludes that points to the *was_client_home*\aspIncludes directory.

Usage notes

The following batch files are provided for client applications to use the ActiveX to EJB bridge:

- **setupCmdLineXJB.bat**
Sets the client environment variables.
- **launchClientXJB.bat**
Calls the setupCmdLineXJB.bat file and launches the application you specify as its arguments; for example:


```
launchClientXJB.bat myapp.exe parm1 parm2
```

or

launchClientXJB MyApplication.vbp

- **Active Server Pages (ASP) include file**

An include file is provided for ASP users to automatically set the following page-level (local) environment variables:

- **com_ibm_websphere_javahome.** Path to the Java run-time directory installed with the WebSphere advanced server client.
- **com_ibm_websphere_washome.** Path to the WebSphere advanced server client directory.
- **com_ibm_websphere_namingfactory.** Sets the Java java.naming.factory.initial system property.
- **com_ibm_websphere_computername.** (Optional) Name of the computer where the WebSphere Advanced Server Client is installed. If you intend to talk to a single specific computer, you are recommended to change this value to become the server name that you intend to access.

- **System settings**

To enable the ActiveX to EJB bridge to access the Java run-time dynamic link library (DLL), the following directories must exist in the system PATH environment variable:

```
was_client_home\java\jre\bin;was_client_home\java\jre\bin\classic
```

Where *was_client_home* is the name of the directory where you installed the WebSphere Application Server client (for example, C:\WebSphere\AppClient).

Note: This technique enables only one Java run time to activate on a machine, therefore all client services on that machine must use the same Java run time. Client applications do not have this limitation because they each have their own private, non-system scope.

JClassProxy and JObjectProxy classes

This topic provides reference information about the object classes of the ActiveX to Enterprise Java Beans (EJB) bridge.

JClassFactory is the object used to access the majority of Java Virtual Machine (JVM) features. This object handles JVM initialization, accesses classes and creates class instances (objects). The majority of tasks for accessing your Java classes and objects are handled with the JClassProxy and JObjectProxy objects:

- XJBInit(String astrJavaParameterArray())

Initializes the JVM environment using an array of strings that represent the command line parameters you normally send to the `java.exe` file.

If you have invalid parameters in the XJBInit() string array, the following error is displayed:

```
Error: 0x6002 "XJBJNI::Init() Failed to create VM" when calling XJBInit()
```

If you have C++ logging enabled, the activity log displays the invalid parameter.

- JClassProxy FindClass(String strClassName)

Uses the current thread class loader to load the specified fully qualified class name and returns a JClassProxy object representing the Java Class object.

- JObjectProxy NewInstance()

Creates a Class instance for the specified JClassProxy object using the parameters supplied to call the Class constructor. For more information about using the JMethodArgs method, see ActiveX to EJB bridge, calling Java methods.

```
JObjectProxy NewInstance(JClassFactory obj, Variant vArg1, Variant vArg2, Variant vArg3, ...)
```

```
JObjectProxy NewInstance(JClassFactory obj, JMethodArgs args)
```

- JMethodArgs GetArgsContainer()

Returns a JMethodArgs object (Class instance).

You can create a JClassProxy object from the JClassFactory.FindClass() method and from any Java method call that normally return a Java Class object. You can use this object as if you had direct access to the Java Class object. All of the class static methods and fields are accessible as are the

java.lang.Class methods. In case of a clash between static method names of the reflected user class and those of the java.lang.Class (for example, getName()), the reflected static methods would execute first.

For example, the following is a static method called getName(). The java.lang.Class object also has a method called getName():

– In Java:

```
class foo{
    foo();
    public static String getName(){return "abcdef";}
    public static String getName2(){return "ghijkl";}
    public String toString2(){return "xyz";}
}
```

– In Visual Basic:

```
...
Dim clsFoo as Object
set clsFoo = oXJB.FindClass("foo")
clsFoo.getName() ' Returns "abcdef" from the static foo class
clsFoo.getName2() ' Returns "ghijkl" from the static foo class
clsFoo.toString() ' Returns "class foo" from the java.lang.Class object.
oFoo = oXJB.NewInstance(clsFoo)
oFoo.toString() ' Returns some text from the java.lang.Object's
                 ' toString() method which foo inherits from.
oFoo.toString2() ' Returns "xyz" from the foo class instance
```

You can create a JObjectProxy object from the JClassFactory.NewInstance() method, and can be created from any Java method call that normally returns a Class instance object. You can use this object as if you had direct access to the Java object and can access all the static methods and fields of the object. All of object instance methods and fields are accessible (including those accessible through inheritance).

The JMethodArgs object is created from the JClassFactory.GetArgsContainer() method. Use this object as a container for method and constructor arguments. You must use this object when overriding the object type when calling a method (for example, when sending a java.lang.String JProxyObject type to a constructor that normally takes a java.lang.Object type).

You can use two groups of methods to add arguments to the collection: Add and Set. You can use Add to add arguments in the order that they are declared. Alternatively, you can use Set to set an argument based on its position in the argument list (where the first argument is in position 1).

For example, if you had a Java Object Foo that took a constructor of Foo (int, String, Object), you could use a JMethodArgs object as shown in the following code extract:

```
...
Dim oArgs as Object
set oArgs = oXJB.GetArgsContainer()

oArgs.AddInt(CLng(12345))
oArgs.AddString("Apples")
oArgs.AddObject("java.lang.Object", oSomeJObjectProxy)

Dim clsFoo as Object
Dim oFoo as Object
set clsFoo = oXJB.FindClass("com.mypackage.foo")
set oFoo = oXJB.NewInstance(clsFoo, oArgs)

' To reuse the oArgs object, just clear it and use the add method
' again, or alternatively, use the Set method to reset the parameters
' Here, we will use Set
oArgs.SetInt(1, CLng(22222))
oArgs.SetString(2, "Bananas")
oArgs.SetObject(3, "java.lang.Object", oSomeOtherJObjectProxy)

Dim oFoo2 as Object
set oFoo2 = oXJB.NewInstance(clsFoo, oArgs)
```

- AddObject (String strObjectName, Object oArg)

Adds an arbitrary object to the argument container in the next available position, casting the object to the class name specified in the first parameter. Arrays are specified using the traditional [] syntax; for example:

```
AddObject("java.lang.Object[][]", oMy2DArrayOfFooObjects)
```

or

```
AddObject("int[]", oMyArrayOfInts)
```

- **AddByte** (Byte byteArg)

Adds a primitive byte value to the argument container in the next available position.

- **AddBoolean** (Boolean bArg)

Adds a primitive boolean value to the argument container in the next available position.

- **AddShort** (Integer iArg)

Adds a primitive short value to the argument container in the next available position.

- **AddInt** (Long lArg)

Adds a primitive int value to the argument container in the next available position.

- **AddLong** (Currency cyArg)

Adds a primitive long value to the argument container in the next available position.

- **AddFloat** (Single fArg)

Adds a primitive float value to the argument container in the next available position.

- **AddDouble** (Double dArg)

Adds a primitive double value to the argument container in the next available position.

- **AddChar** (String strArg)

Adds a primitive char value to the argument container in the next available position.

- **AddString** (String strArg)

Adds the argument in string form to the argument container in the next available position.

- **SetObject** (Integer iArgPosition, String strObjectTypeName, Object oArg)

Adds an arbitrary object to the argument container in the specified position casting it to the class name or primitive type name specified in the second parameter. Arrays are specified using the traditional [] syntax; for example:

```
SetObject(1, "java.lang.Object[][]", oMy2DArrayOfFooObjects)
```

or

```
SetObject(2, "int[]", MyArrayOfInts)
```

- **SetByte** (Integer iArgPosition, Byte byteArg)

Sets a primitive byte value to the argument container in the position specified.

- **SetBoolean** (Integer iArgPosition, Boolean bArg)

Sets a primitive boolean value to the argument container in the position specified.

- **SetShort** (Integer iArgPosition, Integer iArg)

Sets a primitive short value to the argument container in the position specified.

- **SetInt** (Integer iArgPosition, Long lArg)

Sets a primitive int value to the argument container in the position specified.

- **SetLong** (Integer iArgPosition, Currency cyArg)

Sets a primitive long value to the argument container in the position specified.

- **SetFloat** (Integer iArgPosition, Single fArg)

Sets a primitive float value to the argument container in the position specified.

- **SetDouble** (Integer iArgPosition, Double dArg)

Sets a primitive double value to the argument container in the position specified.

- **SetChar** (Integer iArgPosition, String strArg)

Sets a primitive char value to the argument container in the position specified.

- **SetString** (Integer iArgPosition, String strArg)

Sets a java.lang.String value to the argument container in the position specified.

- **Object Item**(Integer iArgPosition)

Returns the value of an argument at a specific argument position.

- Clear()

Removes all arguments from the container and resets the next available position to one.

- Long Count()

Returns the number of arguments in the container.

Java virtual machine initialization tips

Initialize the Java virtual machine (JVM) code with the ActiveX to Enterprise Java Beans (EJB) bridge. For an ActiveX client program (Visual Basic, VBScript, or ASP) to access Java classes or objects, the first step that the program must do is to create Java virtual machine (JVM) code within its process. To create JVM code, the ActiveX program calls the XJBInit() method of the XJB.JClassFactory object. When an XJB.JClassFactory object is created and the XJBInit() method called, the JVM is initialized and ready to use.

- To enable the XJB.JClassFactory to find the Java run-time description definition language (DLL) when initializing, the Java Runtime Environment (JRE) bin and bin\classic directories must exist in the system path environment variable.
- The XJBInit() method accepts only one parameter: an array of strings. Each string in the array represents a command line argument that for a Java program you would normally specify on the Java.exe command line. This string interface is used to set the classpath, stack size, heap size and debug settings. You can get a listing of these parameters by typing java -? from the command line.
- If you set a parameter incorrectly, you receive a 0x6002 "Failed to initialize VM" error message.
- Due to the current limitations of Java Native Interface (JNI), you cannot unload or reinitialize the JVM code after it has loaded. Therefore, after the XJBInit() method has been called once, subsequent calls have no effect other than to create a duplicate JClassFactory object for you to access. It is best to store your XJB.JClassFactory object globally and continue to reuse that object.
- The following Visual Basic extract shows an example of initializing JVM code:

```
Dim oXJB as Object
set oXJB = CreateObject("XJB.JClassFactory")
Dim astrJavaInitProps(0) as String
astrJavaInitProps(0) = _
    "-Djava.class.path=.;c:\myjavaclasses;c:\myjars\myjar.jar"
oXJB.XJBInit(astrJavaInitProps)
```

Example: Developing an ActiveX application client to Enterprise Bean bridge using Java proxy objects

To use Java proxy objects with the ActiveX to Enterprise Java Beans (EJB) bridge:

- After an ActiveX client program (Visual Basic, VBScript, or Active Server Pages (ASP)) has initialized the XJB.JClassFactory object (and thereby, the Java virtual machine (JVM)), it can access Java classes and initialize Java objects. To complete this action, the client program uses the XJB.JClassFactory FindClass() and NewInstance() methods.
- In Java programming, two ways exists to access Java classes: direct invocation through the Java compiler and through the Java Reflection interface. Because the ActiveX to Java bridge needs no compilation and is a complete run-time interface to the Java code, the bridge depends on the latter Reflection interface to access its classes, objects, methods and fields. The XJB.JClassFactory FindClass() and NewInstance() methods behave very similarly to the Java Class.forName() and the Method.invoke() and Field.invoke() methods.
- XJB.JClassFactory.FindClass() takes the fully qualified class name as its only parameter and returns a Proxy Object (JClassProxy). You can use the returned Proxy object like a normal Java Class object and call static methods and access static fields. You can also create a Class Instance (or object), as described below. For example, the following Visual Basic code extract returns a Proxy object for the java.lang.Integer Java class:

```
...
Dim clsMyString as Object
Set clsMyString = oXJB.FindClass("java.lang.Integer")
```

- After the proxy is created, you can access its static information directly. For example, you can use the following code extract to convert a decimal integer to its hexadecimal representation:

```
...
Dim strHexValue as String
strHexValue = clsMyString.toHexString(CLng(255))
```

- The equivalent Java syntax is: `static String toHexString(int i)`. Because ints in Java programming are really 32-bits (which translates to Long in VB), the `CLng()` function converts the value from the default int to a long. Also, even though the `toHexString()` function returns a `java.lang.String`, the code extract does not return an Object proxy. Instead, the returned `java.lang.String` is automatically converted to a native Visual Basic string.

To create an object from a class, you use the `JClassFactory.NewInstance()` method. This method creates an Object instance and takes whatever parameters your class constructor needs. Once the object is created, you have access to all of its public instance methods and fields. For example, you can use the following Visual Basic code extract to create an instance of the `java.lang.Integer` string:

```
...
Dim oMyInteger as Object
set oMyInteger = oXJB.NewInstance(CLng(255))
```

```
Dim strMyInteger as String
strMyInteger = oMyInteger.toString
```

Example: Calling Java methods in the ActiveX to enterprise bean bridge

In the ActiveX to Enterprise Java Beans (EJB) bridge, methods are called using the native language method invocation syntax.

- The following differences between Java invocation and ActiveX Automation invocation exist:
 - Unlike Java methods, ActiveX does not support method (and constructor) polymorphism; that is, you cannot have two methods in the same class with the same name.
 - Java methods are case-sensitive, but ActiveX Automation is not case-sensitive.
- Take care when invoking Java methods through ActiveX Automation. If you use the wrong case on a method call or use the wrong parameter type, you get an Automation Error 438 "Object doesn't support this property or method" thrown.
- To compensate for Java polymorphic behavior, give the exact parameter types to the method call. The parameter types determine the correct method to invoke. For a listing of correct types to use, see ActiveX to EJB bridge, converting data types.
- For example, the following Visual Basic code fails if the `CLng()` method was not present or the `toHexString` syntax was incorrectly typed as `ToHexString`:

```
...
Dim strHexValue as String
strHexValue = clsMyString.toHexString(CLng(255))
```

- Sometimes it is difficult to force some development environments to leave the case of your method calls unchanged. For example, in Visual Basic if you want to call a method `close()` (uncapitalized), the Visual Basic code capitalizes it "`Close()`". In Visual Basic, the only way to effectively work around this behavior is to use the `CallByName()` method. For example:

```
o.Close(123) 'Incorrect...
CallByName(o, "close", vbMethod, 123) 'Correct...
```

or in VBScript, use the `Eval` function:

```
o.Close(123) 'Incorrect...
Eval("o.Close(123)") 'Correct...
```

- The return value of a function is always converted dynamically to the correct type. However, you must take care to use the `set` keyword in Visual Basic. If you expect a non-primitive data type to return, you must use `set`. (If you expect a primitive data type to return, you do not need to use `set`.) For example:

```
Set oMyObject = o.getObject
iMyInt = o.getInt
```

- In some cases, you might not know the type of object returning from a method call, because wrapper classes are converted automatically to primitives (for example, `java.lang.Integer` returns an ActiveX

Automation Long). In such cases, you might need to use your language built-in exception handling techniques to try to coerce the returned type (for example, On Error and Err.Number in Visual Basic).

- **Methods with character arguments**

Because ActiveX Automation does not natively support character types supported by Java methods, the ActiveX to EJB bridge uses strings (byte or VT_11 do not work because characters have multiple bytes in Java code). If you try to call a method that takes a char or java.lang.Character type you must use the JMethodArgs argument container to pass character values to methods or constructors. For more information about how this argument container is used, see *Methods with "Object" Type as Argument and Abstract Arguments*.

- **Methods with "Object" Type as Argument and Abstract Arguments**

Because of the polymorphic nature of Java programming, the ActiveX to Java bridge uses direct argument type mapping to find a method. This method works well in most cases, but sometimes methods are declared with a Parent or Abstract class as an argument type (for example, java.lang.Object). You need the ability to send an object of arbitrary type to a method. To acquire this ability, you must use the XJB.JMethodArgs object to coerce your parameters to match the parameters on your method. You can get a JMethodArgs instance by using the JClassFactory.GetArgsContainer() method.

The JMethodArgs object is a container for method parameters or arguments. This container enables you to add parameters to it one-by-one and then you can send the JMethodArgs object to your method call. The JClassProxy and JObjectProxy objects recognize the JMethodArgs object and attempt to find the correct method and let the Java language coerce your parameters appropriately.

For example, to add an element to a Hashtable object the method syntax is Object put(Object key, Object value). In Visual Basic, the method usage looks like the following example code:

```
Dim oMyHashtable as Object
Set oMyHashtable = _
    oXJB.NewInstance(oXJB.FindClass("java.util.Hashtable"))

' This line will not work. The ActiveX to EJB bridge cannot find a method
' called "put" that has a short and String as a parameter:
oMyHashtable.put 100, "Dogs"
oMyHashtable.put 200, "Cats"

' You must use a XJB.JMethodArgs object instead:
Dim oMyHashtableArgs as Object
Set oMyHashtableArgs = oXJB.GetArgsContainer
oMyHashtableArgs.AddObject("java.lang.Object", 100)
oMyHashtableArgs.AddObject("java.lang.Object", "Dogs")

oMyHashtable.put oMyHashTableArgs
' Reuse the same JMethodArgs object by clearing it.
oMyHashtableArgs.Clear
oMyHashtableArgs.AddObject("java.lang.Object", 200)
oMyHashtableArgs.AddObject("java.lang.Object", "Cats")

oMyHashtable.put oMyHashTableArgs
```

Java field programming tips

Using the ActiveX to Enterprise Java Beans (EJB) bridge to access Java fields has the same case sensitivity issue that it has when invoking methods. Field names must use the same case as the Java field syntax.

- Visual Basic code has the same problem with unsolicited case changing on fields as it does with methods. (For more information about this problem, see *ActiveX to EJB bridge, calling Java methods*). You might need to use the CallByName() function to set a field in the same way that you call a method in some cases. For Fields, you use VBLet for primitive types and VBSet for objects. For example:

```
o.MyField = 123 'Incorrect...
CallByName(o, "MyField", vbLet, 123) 'Correct...
or in VBScript:
```

```
o.MyField = 123
Eval("o.myField = 123")
```

```
'Incorrect...
'Correct...
```

ActiveX to Java primitive data type conversion values

All primitive Java data types are automatically converted to native ActiveX Automation types. However, not all Automation data types are converted to Java types (for example, VT_DATE). Variant data types are used for data conversion. Variant data types are a requirement of any Automation interface, and are used automatically by Visual Basic and VBScript. The tables below provide details about how primitive data types are converted between Automation types and Java types.

Table 4. ActiveX to Java primitive data type conversion

Visual Basic Type	Variant Type	Java Type	Notes
Byte	VT_I1	byte	Byte in Visual Basic is unsigned, but is signed in Java data type.
Boolean	VT_BOOL	boolean	
Integer	VT_I2	short	
Long	VT_I4	int	
Currency	VT_CY	long	
Single	VT_R4	float	
Double	VT_R8	double	
String	VT_BSTR	java.lang.String	
String	VT_BSTR	char	
Date	VT_DATE	n/a	

Example: Using helper methods for data type conversion

Generally, data type conversion between ActiveX (Visual Basic and VBScript) and Java methods occurs automatically, as described in ActiveX to EJB bridge, converting data types. However, the following helper functions are provided for cases where automatic conversion is not possible:

- Byte helper function
- Currency helper function
- Byte helper function

Because the Java Byte data type is signed (-127 through 128) and the Visual Basic Byte data type is unsigned (0 through 255), convert unsigned Bytes to a Visual Basic Integers, which look like the Java signed byte. To make this conversion, you can use the following helper function:

```
Private Function GetIntFromJavaByte(Byte jByte) as Integer
    GetIntFromJavaByte = (CInt(jByte) + 128) Mod 256 - 128
End Function
```

- Currency helper function

Visual Basic 6.0 cannot properly handle 64-bit integers like Java methods can (as the Long data type). Therefore, Visual Basic uses the Currency type, which is intrinsically a 64-bit data type. The only side effect of using the Currency type (the Variant type VT_CY) is that a decimal point is inserted into the type. To extract and manipulate the 64-bit Long value in Visual Basic, use code like the following example. For more details on this technique for converting Currency data types, see Q189862, "HOWTO: Do 64-bit Arithmetic in VBA", on the Microsoft Knowledge Base.

```
' Currency Helper Types
Private Type MungeCurr
    Value As Currency
End Type
Private Type Munge2Long
    LoValue As Long
    HiValue As Long
```

```

End Type

' Currency Helper Functions
Private Function CurrToText(ByVal Value As Currency) As String
    Dim Temp As String, L As Long
    Temp = Format$(Value, "#.0000")
    L = Len(Temp)
    Temp = Left$(Temp, L - 5) & Right$(Temp, 4)
    Do While Len(Temp) > 1 And Left$(Temp, 1) = "0"
        Temp = Mid$(Temp, 2)
    Loop
    Do While Len(Temp) > 2 And Left$(Temp, 2) = "-0"
        Temp = "-" & Mid$(Temp, 3)
    Loop
    CurrToText = Temp
End Function

Private Function TextToCurr(ByVal Value As String) As Currency
    Dim L As Long, Negative As Boolean
    Value = Trim$(Value)
    If Left$(Value, 1) = "-" Then
        Negative = True
        Value = Mid$(Value, 2)
    End If
    L = Len(Value)
    If L < 4 Then
        TextToCurr = CCur(IIf(Negative, "-0.", "0.") & _
            Right$("0000" & Value, 4))
    Else
        TextToCurr = CCur(IIf(Negative, "-", "") & _
            Left$(Value, L - 4) & "." & Right$(Value, 4))
    End If
End Function

' Java Long as Currency Usage Example
Dim LC As MungeCurr
Dim L2 As Munge2Long

' Assign a Currency Value (really a Java Long)
' to the MungeCurr type variable
LC.Value = cyTestIn

' Coerce the value to the Munge2Long type variable
LSet L2 = LC

' Perform some operation on the value, now that we
' have it available in two 32-bit chunks
L2.LoValue = L2.LoValue + 1

' Coerce the Munge value back into a currency value
LSet LC = L2
cyTestIn = LC.Value

```

Array tips for ActiveX application clients

Arrays are very similar between Java and Automation containers like Visual Basic and VBScript. Here are some important points to consider when passing arrays back and forth between these containers:

- Java arrays cannot mix types. All Java arrays contain a single type, so when passing arrays of variants to a Java array, you must make sure that all of the elements in the variant array are of the same base type. For example, in Visual Basic code:

```

...
Dim VariantArray(1) as Variant
VariantArray(0) = CLng(123)
VariantArray(1) = CDb1(123.4)
oMyJavaObject.foo(VariantArray) ' Illegal!

```

```
VariantArray(0) = CLng(123)
VariantArray(1) = CLng(1234)
oMyJavaObject.foo(VariantArray) ' This works
```

- Arrays of primitive types are converted using the rules defined in primitive data type conversion.
- Arrays of Java objects are handled through arrays of JObjectProxy objects.
- Arrays of JObjectProxy objects must be fully initialized and of the correct associated Java type. When initializing an array in Visual Basic (for example, Dim oJavaObjects(1) as Object), you must set each object to a JObjectProxy object before you send the array to a Java object. The bridge is unable to determine the type of null or empty object values.
- When receiving an array from a Java method, the lower-bound is always zero. Java methods only support zero-based arrays.
- Nested or multidimensional arrays are treated as zero-based multidimensional arrays in Visual Basic and VBScript containers.
- Uninitialized arrays or Array Types are unsupported. When calling a Java method that takes an array of objects as a parameter, you must fully initialize the array of JObjectProxy objects.

Error handling codes for ActiveX application clients

All exceptions thrown in Java code are encapsulated and thrown again as a COM error through the ISupportErrorInfo interface and the EXCEPINFO structure of IDispatch::Invoke(); the Err object in Visual Basic and VBScript. Because there are no error numbers associated with Java exceptions, whenever a Java exception is thrown, the entire stack trace is stored in the error description text and the error number assigned is 0x6003.

In Visual Basic or VBScript, you need to use the Err.Number and Err.Description fields to determine the actual Java error. Non-Java errors are thrown as you would expect via the IDispatch interface; for example, if a method cannot be found, then error 438 "Object doesn't support this property or method" is thrown.

Error number	Description
0x6001	Java Native Interface (JNI) error
0x6002	Initialization error
0x6003	Java exception. Error description is the Java Stack Trace.
0x6FFF	General Internal Failure

Threading tips

The ActiveX to Enterprise Java Beans (EJB) bridge supports both free-threaded and apartment-threaded access and implements the Free Threaded Marshaler to work in a hybrid environment such as Active Server Pages (ASP). Each thread created in the ActiveX process is mirrored in the Java environment when the thread communicates through the ActiveX to EJB bridge. In addition, once all references to Java objects (there are no JObjectProxy or JClassProxy objects) are loaded in an ActiveX thread, the ActiveX to EJB bridge detaches the thread from the Java virtual machine (JVM) code. Therefore, you must be careful that any Java code that you access from a multithreaded Windows application is thread safe. Visual Basic code and VBScript applications are both essentially single threaded. Therefore, Visual Basic and VBScript applications do not have threading issues in the Java programs they access. Active Server Pages and multithreaded C and C++ programs can have issues.

Consider the following scenario:

1. A multithreaded Windows Automation Container (our ActiveX Process) starts. It exists on Thread A.
2. The ActiveX Process initializes the ActiveX to EJB bridge, which starts the JVM code. The JVM attaches to the same thread and internally calls it Thread 1.
3. The ActiveX Process starts two threads: B and C.
4. Thread B in the ActiveX Process uses the ActiveX to EJB bridge to access an object that was created in Thread A. The JVM attaches to thread B and calls it Thread 2.

5. Thread C in the ActiveX Process never talks to the JVM code, so the JVM never needs to attach to it. This is a case where the JVM code does not have a one-to-one relationship between ActiveX threads and Java threads.
6. Thread B later releases all of the JObjectProxy and JClassProxy objects that it used. The Java Thread 2 is detached.
7. Thread B again uses the ActiveX to EJB bridge to access an object that was created in Thread A. The JVM code attaches again to the thread and calls it Thread 3.

ActiveX process	JVM access by ActiveX process
Thread A - Created in 1	Thread 1 - Attached in 2
Thread B - Created in 4	Thread 2 - Attached in 4, detached in 6 Thread 3 - Attached in 7
Thread C - Created in 4	

Threads and Active Server Pages

Active Server Pages (ASP) in Microsoft Internet Information Server is a multithreaded environment. When you create the XJB.JClassFactory object, you can store it in the Application collection as an Application-global object. All threads within your ASP environment can now access the same ActiveX to EJB bridge object. Active Server Pages by default creates 10 Apartment Threads per ASP process per CPU. This means that when your ActiveX to EJB bridge object is initialized any of the 10 threads can call this object, not just the thread that created it.

If you need to simulate single-apartment behavior, you can create a Single-Apartment Threaded ActiveX dynamic link library (DLL) in Visual Basic code and encapsulate the ActiveX to the EJB bridge object. This encapsulation guarantees that all access to the JVM object is on the same thread. You need to use the <OBJECT> tag to assign the XJB.JClassFactory to an Application object and must be aware of the consequences of introducing single-threaded behavior to a Web application.

The Microsoft KnowledgeBase has several articles about ASP and threads, including:

- Q243543 INFO: Do Not Store STA Objects in Session or Application
- Q243544 INFO: Component Threading Model Summary Under Active Server Pages
- Q243548 INFO: Design Guidelines for VB Components Under ASP

Example: Viewing a System.out message

The ActiveX to Enterprise Java Beans (EJB) bridge does not have a console available to view Java System.out messages. To view these messages when running a stand-alone client program (such as Visual Basic), you need to redirect the output to a file. For example:

```
launchClientXJB.bat MyProgram.exe > output.txt
```

- To view the System.out messages when running a Service program such as Active Server Pages, you need to override the Java System.out OutputStream object to FileOutputStream. For example, in VBScript:

```
'Redirect system.out to a file
' Assume that oXJB is an initialized XJB.JClassFactory object
Dim clsSystem
Dim oOS
Dim oPS
Dim oArgs

' Get the System class
Set clsSystem = oXJB.FindClass("java.lang.System")

' Create a FileOutputStream object
' Create a PrintStream object and assign to it our FileOutputStream
Set oArgs = oXJB.GetArgsContainer oArgs.AddObject "java.io.OutputStream", oOS
```



```
Set oPS = oXJB.NewInstance(oXJB.FindClass("java.io.PrintStream"), oArgs)
```

```
' Set our System OutputStream to our file  
clsSystem.setOut oPS
```

Example: Enabling logging and tracing for application clients

The ActiveX to EJB bridge provides two logging and tracing formats: Windows Application Event Log and Java Trace Log.

- Windows Event Log

The Windows Application Event Log shows JNI errors, Java console error messages, and XJB initialization messages. This log is most useful for determining XJBInit() errors and any unusual exceptions that do not come from the Java environment. By default, critical error logging will be enabled and debug and event logging is disabled.

To enable or disable logging of certain event types to the Windows Event Log, specify one or more parameters to XJBInit(). If more than one parameter is set, they will be processed in the order in which they appear in the input string array to the XJBInit() method. Once the XJBInit() method is initialized, these parameters can no longer be set/reset for the life of the process. Using Java java.lang.System.setProperty() to set these values also has no effect.

- -Dcom.ibm.ws.client.xjb.native.logging.debug=enabled|disabled

Enables or disables debug level messages from displaying in the Windows operating system event log. This level of logging is most useful and shows most internal errors, user programming issues or configuration problems.

- -Dcom.ibm.ws.client.xjb.native.logging.event=enabled|disabled

Enables or disables event level messages from appearing in the Windows operating system event log.

- -Dcom.ibm.ws.client.xjb.native.logging.*=enabled|disabled

Enables or disables both event and debug level messages from appearing in the Windows operating system event log. It is not possible to disable some critical error messages from being displayed in the error log. Only debug and event level messages can be disabled.

Viewing the Windows application event log with the event viewer:

To open the event viewer in the Windows operating system:

1. Click **Start > Settings > Control Panel**.
2. Double-click **Administrative Tools**.
3. Double-click **Event Viewer**.

All ActiveX to EJB bridge events display the text WebSphere XJB in the source column and in the application log. For information about using Event Viewer, click the **Action** menu in Event Viewer, and then click **Help**.

To open the even viewer in the Windows operating system, click **Start > Programs > Administrative Tools > Event Viewer**. All ActiveX to EJB bridge events have the text WebSphere XJB in the source column and display in the application log. For information about using Event Viewer, click the **Help** menu in Event Viewer.

- Java trace log

The Java trace log displays information that you can use to debug method calls, class lookups, and argument coercion problems. Since the Java portion of the bridge mirrors the function of the COM IDispatch interface, the information in the trace log is similar to what you have come to expect from an IDispatch interface. To understand the trace log, you need a fundamental understanding of IDispatch.

To enable user-logging, add the following parameters to the XJBInit() input string array:

```
"-DtraceString=com.ibm.ws.client.xjb.*=event=enabled"  
"-DtraceFile=C:\MyTrace.txt"
```

ActiveX client programming best practices

In general, the best way to access Java components is to use the Java language. It is recommended that you do as much programming as possible in the Java language and use a small simple interface between

your COM Automation container (for example, Visual Basic) and the Java code. This interface avoids any overhead and performance problems that can occur when moving across the interface.

- Visual Basic guidelines
- Active Server Pages guidelines
- J2EE guidelines

Visual Basic guidelines

The following guidelines are intended to help optimize your use of the ActiveX to EJB bridge with Visual Basic:

- Launch the Visual Basic replication through the `launchClientXJB.bat` file. If you want to run your Visual Basic application through the Visual Basic debugger, run the Visual Basic integrated development environment (IDE) within the ActiveX to EJB bridge environment. After you create your Visual Basic project, you can launch it from a command line; for example, `launchClientXJB MyApplication.vbp`. You can also launch the Visual Basic application alone in the ActiveX to EJB environment, by changing the Visual Basic shortcut on the Windows Start menu so that the `launchClientXJB.bat` file precedes the call to the `VB6.EXE` file.
- Exit the Visual Basic IDE before debugging programs.

Because the Java virtual machine (JVM) code attaches to the running process, you must exit the Visual Basic editor before debugging your program. If you run the process, then exit your program within the Visual Basic IDE, the JVM code continues to run and you reattach the same JVM code when `XJBInit()` is called by the debugger. This causes problems if you try to update `XJBInit()` arguments (for example, classpath) because the changes are not be applied until you restart the Visual Basic program.

- Store the `XJB.JClassFactory` object globally.

Because you cannot unload or reinitialize the JVM code, cache the resulting `XJB.JClassFactory` object as a global variable. The overhead of treating this object as a global variable or passing a single reference around is much less than recreating a new `XJB.JClassFactory` object and calling the `XJBInit()` argument more than once.

CScript and Windows Scripting Host

The following guidelines intend to help optimize your use of the ActiveX to EJB bridge with CScript and Windows Scripting Host (WSH):

- Launch in ActiveX to EJB environment.

Launch the VBScript files in the ActiveX to EJB bridge environment, to run VBScript files in `.vbs` files.

Two common ways exist to launch your script:

- `launchClientXJB MyScript.vbs`
- `launchClientXJB cscript MyScript.vbs`

Active Server Pages guidelines

The following guidelines intend to help optimize your use of the ActiveX to EJB bridge with Active Server Pages software:

- Use the ActiveX to EJB Helper functions from the Active Server Pages Application.

Because Active Server Pages (ASP) code typically use VBScript, you can use the included helper functions in any VBScript environment with minor changes. For more information about these helper functions, see *Helper functions for data type conversion*. To run outside of the ASP environment, remove or change all references to the `Server`, `Request`, `Response`, `Application` and `Session` objects; for example, change `Server.CreateObject` to `CreateObject`.

- Set JRE path globally in system.

The `XJB.JClassFactory` object must be able to find the Java run time dynamic link library (DLL) when initializing. In Internet Information Server, you cannot specify a path for its processes independently; you must set the process paths in the system `PATH` variable. You can only have a single JVM version

available on a machine using the ASP application. Also, remember that after you change the system PATH variable you must reboot the Internet Information Server machine so that the Internet Information Server can see the change.

- Set the system TEMP environment variable.

If the system TEMP environment variable is not set, Internet Information Server stores all temporary files in the WINNT directory, which is usually not desired.

- Use high isolation or an isolated process.

When using the ActiveX to Java bridge with Active Server Pages software, creating your Web application in its own process is recommended. You can only load one JVM instruction in a single process and if you want to have more than one application running with different JVM environment options (for example, different classpaths), then you need to have separate processes.

- Use the Application Unload option.

When debugging your application, use **Unload** when viewing your ASP application properties in the Internet Information Server administration console to unload the process from memory and thereby unload the JVM code.

- Run one process per application.

Use only one ASP application per J2EE application or JVM environment, in your ASP environment. If you need separate class paths or JVM settings, you need separate ASP applications (virtual directories with high isolation or an isolated process).

- Store the XJB.JClassFactory object in application scope.

Because of the one-to-one relationship required between a JVM instruction and a process, and because the JVM code can never detach or shut down from a process independently, cache the XJB.JClassFactory object at application scope and call the XJBInit() method only once.

Because the ActiveX to EJB bridge employs a free-threaded marshaler, take advantage of the multi-threaded nature of Internet Information Server and the ASP environment. If you choose to reinitialize the XJB.JClassFactory object at Page scope (local variables), then the XJBInit() method can only initialize your local XJB.JClassFactory variable. It is more efficient to use the XJBInit() method once.

- Use VBScript conversion functions.

Because VBScript code only supports variant data types, use the CStr(), CByte(), CBool(), CCur(), CInt(), CInq(), CSng() and CDbl() functions to tell the activeX to EJB bridge which data type you are using; for example `oMyObject.Foo(CDb1(1.234))`.

J2EE guidelines

The following guidelines are intended to help optimize your use of the ActiveX to EJB bridge with the J2EE environment;

- Store client container objects globally.

Because you can only have one JVM instruction per process, and a single J2EE client container (`com.ibm.websphere.client.applicationclient.launchClient`) per JVM instruction, initialize your J2EE client container only once and reuse it. For ASP applications, store the J2EE client container in an application level variable and initialize it only once (either on the `Application_OnStart()` event in the `global.asa` file or by checking to see if it `IsEmpty()`).

A side effect to storing the client container object globally is that you cannot change the client container parameters without destroying the object and creating a new one. These parameters include the EAR file, BootstrapHost, class path, and so on. If you run a Visual Basic application and want to change the client container parameters, you must end the application and restart it. If you run an Active Server Pages application, you must first unload the application from Internet Information Server (see "Use the Application Unload Button" under Active Server Pages guidelines). Then load the Active Server Pages application with the different client container parameters. The parameters set the first time the Active Server Pages application loads. Since the client container is stored on the Internet Information Server, all the browser clients share the parameters using the Active Server Pages application. This behavior is normal for Active Server Pages code, but can be confusing when you try to run to different WebSphere Application Servers using the same Active Server Pages application, which is not supported.

- Reuse custom temporary directory for EAR file extraction.

By default, the client container launches and extracts the application .ear file to your temp directory and then sets up the thread class loader to use the extracted EAR file directory and the JAR files included in the client JAR manifest. This process is time consuming and because of some limitations with JVM shutdown through Java Native Interface (JNI) and file locking, these files are never cleaned up.

Specifically, each time the client container launch() method is called, it extracts the EAR file to a random directory name in your temporary directory on your hard drive. The current Java thread class loader is then changed to point to this extracted directory which in turn locks the files within. In a normal J2EE Java client, these files automatically clean up after the application exits. This cleanup occurs when the client container shutdown hook is called (which never happens in the ActiveX to EJB bridge), which leaves the temporary directory there.

To avoid these problems, you can specify a directory to extract the EAR file by setting the com.ibm.websphere.client.applicationclient.archivedir Java system property before calling the client container launch() method. If the directory does not exist or is empty, you extract the EAR file normally. If the EAR file was previously extracted, the directory is reused. This feature is particularly important for server processes (for example, ASP), which can stop and restart, potentially calling the launchClient() method several times.

If you need to update your EAR file, delete the temporary directory first. The next time you create the client container object, it extracts the new EAR file to the temporary directory. If you do not delete the temporary directory or change the system property value to point to a different temporary directory, the client container reuses the currently extracted EAR file, and does not use your changed EAR file.

Note: When specifying the com.ibm.websphere.client.applicationclient.archivedir property, ensure that the directory you specify *is unique* for each EAR file you use. For example, do not point MyEar1.ear and MyEar2.ear files to the same directory.

If you choose not to use this system property, go regularly to your Windows temp directory and delete the WSTMP* subdirectories. Over a relatively short period of time, these subdirectories can waste a significant amount of space on the hard drive.

Developing applet client code

Applet clients have the following setup requirements:

- These clients are available on the Windows XP, Windows NT or Windows 2000 platforms. Check the prerequisites page for information on new platform support.
- They require one of these browsers:
 - Internet Explorer version 5.0 and later
 - Netscape Navigator 4.7 and later
- The browser installation precedes the client code installation.

Unlike typical applets that reside on either Web servers or WebSphere Application Servers and can only communicate using the HTTP protocol, applet clients are capable of communicating over the HTTP protocol and the RMI-IIOP protocol. This additional capability gives the applet direct access to enterprise beans.

1. Run the application server client installation.
2. Select the applet client option.
3. Install an applet client.
4. Install the WebSphere Application Server plug-in for the browser. From the WebSphere Application Server Java Plug-in control panel, enter the following code:

```
-Djava.security.policy=<product_installation_dir>\properties\client.policy
-Dwas.install.root=<product_installation_dir>
-Djava.ext.dirs=<product_installation_dir>\classes;
<product_installation_dir>\java\jre\lib\ext;
```

```
<product_installation_dir>\java\jre\lib;  
<product_installation_dir>\lib;<product_installation_dir>\properties  
-Dcom.ibm.CORBA.ConfigURL=file:<product_installation_dir>\properties\sas.client.props  
-classpath <product_installation_dir>\properties
```

Note: The previously entries are automatically placed into the WebSphere Application Server control panel for the Java plug-in user who installed the WebSphere Application Server Application Client. If this sample is being run by a user other than the person who installed the client, the user must enter the entries.

- The Java **Run-Time Parameters** field is similar to the command prompt when using command line options. Therefore, you can enter most options available from the command prompt (for example, -cp, classpath, and others) in this field as well.
- Access the control panel from the **Start** menu. Click **Start > Control panel >** select the product Java plug-in.
- The applet container is the Web browser and the Java plug-in combination. You must first install the WebSphere Application Server Applet client so that the browser recognizes the IBM product Java plug-in.

View the Samples gallery for more information about application clients. Before you run the basicCalculator Sample, ensure the JMS Server is started. (See "Samples Gallery" in the Information Center.)

These samples do not include client applications that communicate with Enterprise server-side examples.

Accessing secure resources using SSL and applet clients

By default, the applet client is configured to have security enabled. If you have global security turned on at the server from which you are accessing resources, then you can use secure sockets layer (SSL) when needed. If you decide that the security requirements for the applet differ from other application client types, then create a new version of the sas.client.props file.

1. Make a copy of the following file so that you can use it for an applet:
`<product_install_directory>/properties/sas.client.props`
2. Edit the copy of sas.client.props file that you made with your changes.
3. Click **Start > Control panel >** select the product Java plug-in to open the Java control panel.
 - To use the file you created in step 1, modify the following value:

```
-Dcom.ibm.CORBA.ConfigURL=file:<product_install_directory>\properties\sas.client.props
```

For more information on the sas.client.props file and WebSphere Application Server security, see the Security section of the information center.

Applet client security requirements

When code is loaded, it is assigned permissions based on the security policy in effect. This policy specifies the permissions that are available for code from various locations. You can initialize this policy from an external policy file. By default, the client uses the `<product_installation_dir>/properties/client.policy` file. You must update this file with the following permission:

SocketPermission grants permission to open a port and make a connection to a host machine, which is your WebSphere Application Server. In the following example, yourserver.yourcompany.com is the complete host name of your WebSphere Application Server:

```
permission java.util.PropertyPermission "*", "read";  
permission java.net.SocketPermission "yourserver.yourcompany.com", "connect";
```

Applet client tag requirements

Standard applets require the HTML `<APPLET>` tag to identify the applet to the browser. The `<APPLET>` tag invokes the Java virtual machine (JVM) of the browser.

- For applets to communicate with Enterprise Java Beans (EJB) in the WebSphere Application Server environment, the <APPLET> tag must be replaced with the following tags:

```
<OBJECT>
<EMBED>
```

- The classid and type attributes cannot be modified, and must be entered as described in the applet client example. The codebase attribute on the <OBJECT> tag must be excluded. Do not confuse the codebase attribute on the <OBJECT> tag with the codebase attribute on the <PARAM> tag. Although both attributes are called codebase, they are separate entities.
- The following code example illustrates the applet code. In this example, MyApplet.class is the applet code, applet.jar is the file that contains the applet code, and EJB.jar is the file that contains the enterprise bean code:

```
<OBJECT classid="clsid:8AE2D840-EC04-11D4-AC77-006094334AA9"
width="600" height="500">
<PARAM NAME=CODE VALUE=MyAppletClass.class>
<PARAM NAME="archive" VALUE='Applet.jar, EJB.jar'>
<PARAM TYPE="application/x-java-applet;version=1.3">
<PARAM NAME="scriptable" VALUE="false">
<PARAM NAME="cache-option" VALUE="Plugin">
<PARAM NAME="cache-archive" VALUE="Applet.jar, EJB.jar">
<COMMENT>
<EMBED type="application/x-websphere-client" CODE=MyAppletClass.class
ARCHIVE="Applet.jar, EJB.jar" WIDTH="600" HEIGHT="500"
scriptable="false">
<NOEMBED>
</COMMENT>
</NOEMBED>WebSphere Java Application/Applet Thin Client for
Windows is required.
</EMBED>
</OBJECT>
```

- The value of the type attribute on the </EMBED> tag can also be, for example:

```
<EMBED type="application/x-websphere-client, version=4.0" ...
```

Applet client code requirements

The code used by an applet to talk to an enterprise bean is the same as that used by a stand-alone Java program or a servlet, except for one additional property called `java.naming.applet`. This property informs the `InitialContext` and the Object Request Broker (ORB) that this client is an applet rather than a stand-alone Java application or servlet.

- When you initialize an instance of the `InitialContext` class, the first two lines in this code snippet illustrate what both a stand-alone Java program and a servlet issue to specify the computer name, domain, and port. In this example, `<yourserver.yourdomain.com>` is the computer name and domain where WebSphere Application Server resides, and 900 is the configured port. After the bootstrap values (`<yourserver.yourdomain.com>:900`) are defined, the client to server communications occur within the underlying infrastructure. In addition to the first two lines for applets, you must add the highlighted third line to your code. That highlighted line identifies this program as an applet, for example:

```
prop.put(Context.INITIAL_CONTEXT_FACTORY, "com.ibm.websphere.naming.WsnInitialContextFactory");
prop.put(Context.PROVIDER_URL, "iiop://<yourserver.yourdomain.com>:900)
prop.put(Context.APPLET, this);
```

Developing J2EE application client code

A *J2EE application client* program operates similarly to a standard J2EE program in that it runs its own Java virtual machine (JVM) code and is invoked at its main method.

The Java Virtual Machine application client program differs from a standard Java program because it uses the Java Naming and Directory Interface (JNDI) namespace to access resources. In a standard Java program, the resource information is coded in the program.

Storing the resource information separately from the client application program makes the client application program portable and more flexible.

1. Write the client application program. Write the J2EE application client program on any development machine. At this stage, you do not require access to the WebSphere Application Server.

Using the `javax.naming.InitialContext` class, the client application program uses the look-up operation to access the Java Naming and Directory Interface (JNDI) namespace. The `InitialContext` class provides the `lookup` method to locate resources.

The following example illustrates how a client application program uses the `InitialContext` class:

```
import javax.naming.*

public class MyAppClient
{
    public static void main(String argv[])
    {
        InitialContext initCtx = new InitialContext();
        Object homeObject = initCtx.lookup("java:comp/env/ejb/BasicCalculator");
        BasicCalculatorHome bcHome = (BasicCalculatorHome)
        javax.rmi.PortableRemoteObject.narrow(homeObject, BasicCalculatorHome.class);
        BasicCalculatorHome bc = bcHome.create();          ...
    }
}
```

In this example, the program looks up an enterprise bean called `BasicCalculator`. The `BasicCalculator` EJB reference is located in the client JNDI namespace at `java:comp/env/ejb/BasicCalculator`. Since the actual Enterprise Java Bean runs on the server, the application client run time returns a reference to the `BasicCalculator` home interface.

If the client application program lookup was for a resource reference or an environment entry, then the look up function returns an instance of the configured type as defined by the client application deployment descriptor. For example, if the program lookup was a JDBC data source, the lookup would return an instance of `javax.sql.DataSource`. Although you can edit deployment descriptor files, do not use the administrative console to modify them.

2. Assemble the application client using the Assembly Toolkit.

The JNDI namespace knows what to return on a lookup because of the information assembled by the assembly tool.

Assemble the J2EE application client on any development machine with the assembly tool installed.

When you assemble your application client, provide the application client run time with the required information to initialize the execution environment for your client application program.

Remember following when you configure resources used by your client application program:

- Resource environment references are different than resource references. Resource environment references allow your application client to use a logical name to look-up a resource bound into the server JNDI namespace. A resource reference allows your application to use a logical name to look up a local J2EE resource. The J2EE specification does not specify a particular implementation of a resource. The following table contains supported resource types and identifies the resources to which the WebSphere Application Server provides a client implementation.

Resource Type	Client Configuration Notes	Client implementation provided by WebSphere Application Server
<code>javax.sql.DataSource</code>	Supports specification of any data source implementation class	No
<code>java.net.URL</code>	Supports specification of custom protocol handlers	Provided by Java Runtime Environment files
<code>javax.mail.Session</code>	Supports custom protocol configuration	Yes - POP3, SMTP, IMAP

Resource Type	Client Configuration Notes	Client implementation provided by WebSphere Application Server
javax.jms.QueueConnectionFactory, javax.jms.TopicConnectionFactory, javax.jms.Queue, javax.jms.Topic	Supports configuration of WebSphere embedded messaging, IBM MQ Series and other JMS providers	Yes - WebSphere embedded messaging

3. Assemble the Enterprise Archive (EAR) file.

The application is contained in an enterprise archive or .ear file. The .ear file is composed of:

- Enterprise bean, application client, and user-defined modules or .jar files
- Web applications or .war files
- Metadata describing the applications or application .xml files

You must assemble the .ear file on the server machine.

4. Distribute the EAR file.

The client machines configured to run this client must have access to the .ear file.

If all the machines in your environment share the same image and platform, run the Application Client Resource Configuration Tool (ACRCT) on one machine to configure the external resources, and then distribute the configured .ear file to the other machines.

If your environment is set up with a variety of client installations and platforms, run the ACRCT for each unique configuration.

You can either distribute the .ear files to the correct client machines, or make them available on a network drive.

Distributing the .ear files is the responsibility of the system and network administrator.

5. Deploy the application client.

6. Configure the application client resources.

If the client application defines the local resources, run the ACRCT (clientConfig command) on the local machine to reconfigure the .ear file. Use the ACRCT to change the configuration. For example, the .ear file can contain a DB2 resource, configured as C:\DB2. If, however, you installed DB2 in the D:\Program Files\DB2 directory, use the ACRCT to create a local version of the .ear file.

After developing the J2EE application client code, launch the application client.

J2EE application client class loading

When you run your J2EE application client, a hierarchy of class loaders is created to load classes used by your application.

The following list describes the hierarchy of class loaders:

- The topmost class loader, the bootstrap class loader, contains the JAR files that make up the Java virtual machine code, such as rt.jar, plus those JAR files defined by the -Xbootclasspath parameter on the Java command. The WebSphere application client run time sets this value to the WAS_BOOTCLASSPATH environment variable.
- The *extensions class loader* class loader is a child to the bootstrap class loader. This class loader contains JAR files in the java/jre/lib/ext directory or those JAR files defined by the -Djava.ext.dirs parameter on the Java command. The WebSphere application client run time does not set -Djava.ext.dirs parameters, so it uses the JAR files in the java/jre/lib/ext directory.
- The *system class loader* class loader contains JAR files and classes that are defined by the -classpath parameter on the Java command. The WebSphere application client run time sets this parameter to the WAS_CLASSPATH environment variable.
- The *WebSphere class loader* class loader loads the WebSphere application client run time and any classes placed in the WebSphere application client user directories. The directories used by this class loader are defined by the WAS_EXT_DIRS environment variable. The WAS_BOOTCLASSPATH, WAS_CLASSPATH,

and the `WAS_EXT_DIRS` environment variables are set in the `installation_root/bin/setupCmdLine` command shell for WebSphere Application Server installations, or in the `installation_root/bin/setupClient` command shell for client installations.

As the J2EE application client run time initializes, additional class loaders are created as children of the WebSphere class loader. If your client application uses resources such as Java Database Connectivity (JDBC) API, Java Message Service (JMS) API, or Uniform Resource Locator (URL), a different class loader is created to load each of those resources. Finally, the application client run time sets the WebSphere class loader to load classes within the `.ear` file by processing the client JAR manifest repeatedly. The system class path, defined by the `CLASSPATH` environment variable is never used and is not part of the hierarchy of class loaders.

To package your client application correctly, you must understand which class loader loads your classes. When the Java code loads a class, the class loader used to load that class is assigned to it. Any classes subsequently loaded by that class will use that class loader or any of its parents, but it will not use children class loaders.

In some cases the WebSphere application client run time can detect when your client application class is loaded by a different class loader from the one created for it by the WebSphere application client run time. When this detection occurs, you see the following message:

```
WSCL0205W: The incorrect class loader was used to load [0]
```

This message occurs when your client application class is loaded by one of the parent class loaders in the hierarchy. This situation is typically caused by having the same classes in the `.ear` file and on the hard drive. If one of the parent class loaders locates a class, that class loader loads it before the application client run time class loader. In some cases, your client application still functions correctly. In most cases, however, you receive "class not found" exceptions.

Configuring the classpath fields

When packaging your J2EE client application, you must configure various class path fields. Ideally, you should package everything required by your application into your `.ear` file. This is the easiest way to distribute your J2EE client application to your clients. However, you should not package such resources as JDBC APIs, JMS APIs, or URLs. In the case of these resources, use class path references to access those classes on the hard drive. You might also have other classes installed on your client machines that you do not need to redistribute. In this case, you also want to use classpath references to access the classes on the hard drive, as described below.

Referencing classes within the EAR file

WebSphere product J2EE applications do not use the system class path. Use the MANIFEST Class path entry to refer to other JAR files within the `.ear` file. Configure these values using the module Class path fields in the Application Assembly Tool Assembly Toolkit. For example, if your client application needs to access the path of the EJB JAR file, add the deployed enterprise bean module name to your application client Class path field in the Assembly Tool. The format of the Class path field for each of the different modules (Application Client, EJB, Web) is the same:

- The values must refer to `.jar` and `.class` files that are contained within the `.ear` file.
- The values must be relative to the root of the `.ear` file.
- The values cannot refer to absolute paths in the file systems.
- Multiple values must be separated by spaces, not colons or semi-colons.

Note: This is the Java method for allowing applications to function platform independent.

Typically, you add modules (`.jar` files) to the root of the `.ear` file. In this case, you only need to specify the name of the module (`.jar` file) in the Class path field. If you choose to add a module with a path, you need to specify the path relative to the root of the `.ear` file.

For referencing .class files, you must specify the directory relative to the root of the .ear file. With the Assembly Tool you can add individual class files to the .ear file. It is recommended that these additional class files are packaged in a .jar file. Add this .jar file to the module Class path fields. If you add .class files to the root of the .ear file, add ./ to the module Class path fields. Consider the following example directory structure in which the file myapp.ear contains an application client JAR file named client.jar and a mybeans.jar EJB module. Additional classes reside in class1.jar and utility/class2.zip files. A class named xyz.class is not packaged in a JAR file but is in the root of the EAR file. Specify **./ mybeans.jar utility/class2.zip class1.jar** as the value of the Classpath property. The search order is:
myapp.ear/client.jar myapp.ear/xyz.class myapp.ear/mybeans.jar myapp.ear/utility/class2.zip myapp.ear/class1.jar

Referencing classes that are not in the EAR file

Use the launchClient -CCclasspath parameter. This parameter is specified at run time and takes platform-specific class path values, which means multiple values are separated by semi-colons or colons. The client and the server are similar in this respect.

Resource class paths

When you configure resources used by your client application using the Application Client Resource Configuration Tool, you can specify class paths that are required by the resource. For example, if your application is using a JDBC to a DB2 database, add db2java.zip to the class path field of the database provider. These class path values are platform-specific and require semi-colons or colons to separate multiple values.

Using the launchClient API

If you use the launchClient shell and bat command, the WebSphere class loader hierarchy is created for you. However, if you use the launchClient API, you must perform this setup yourself. Copy the launchClient shell command in defining the Java system properties.

Developing pluggable application client code

As you prepare to install the pluggable application client, remember that pluggable clients are only available on Windows systems.

1. Install the pluggable application client from the WebSphere Application Client CD by selecting option **Pluggable Application Client** from the **Custom client installation** panel.
2. Set the Java application pluggable client environment by using the **setupClient** shell, located in:
install_root\AppClient\bin\setupClient.bat
3. Add your specific Java client application JAR files to the CLASSPATH and start your Java client application from this environment, after setting the environment variables.
4. Run the following Java command to invoke your client application:

```
%JAVA_HOME%/bin/java -Xbootclasspath/p:%WAS_BOOTCLASSPATH% -classpath  
<list of your application jars and classes> -Djava.ext.dirs=%WAS_EXT_DIRS%  
-Djava.naming.provider.url=iiop://<your WebSphere server machine name>  
-Djava.naming.factory.initial=com.ibm.websphere.naming.WsnInitialContextFactory  
%SERVER_ROOT% %CLIENTSAS% <fully qualified class name to run>  
  
$JAVA_HOME/bin/java -Xbootclasspath/p:$WAS_BOOTCLASSPATH -classpath  
<list of your application jars and classes> -Djava.ext.dirs=$WAS_EXT_DIRS  
-Djava.naming.provider.url=iiop://<your WebSphere server machine name>  
-Djava.naming.factory.initial=com.ibm.websphere.naming.WsnInitialContextFactory  
$SERVER_ROOT $CLIENTSAS <fully qualified class name to run>
```

View the Samples Gallery for more information about application clients. Before you run the basicCalculator Sample, ensure the JMS Server is started. (See "Samples Gallery" in the Information Center.)

These samples do not include client applications that communicate with Enterprise server-side examples.

Developing thin application client code

You can develop and run Java thin client applications on machines installed with either a client or a server. The client provides a setup command shell which sets up your environment for either a thin client application or a J2EE client application. The server provides a command shell which sets up your environment for J2EE application clients only. The Java invocation to run a thin application client varies between a client and a server. If your thin client application needs to run on both a client installation and a server installation, follow the steps for developing thin application clients on a server machine.

1. Install the Java application thin client from the WebSphere Application Client CD by selecting option **J2EE/Thin application client for the WebSphere Application Client**.
2. Perform one of the following:
 - Develop thin application client code for a client machine.
 - Develop thin application client code for a server machine.

View the Samples gallery for more information about application clients. (See "Samples Gallery" in the Information Center.)

These samples do not include client applications that communicate with Enterprise server-side examples.

Developing thin application client code on a client machine

You must install the thin application client from the WebSphere Application Client CD before performing this task. For more information, see Developing thin application client code.

1. Set the Java application thin client environment by using the setupClient shell, located in: Windows systems:

```
install_root\AppClient\bin\setupClient.bat
```

UNIX platforms:

```
install_root/AppClient/bin/setupClient.sh
```

2. Run the following Java compilation command to compile your client application. On Windows systems, enter:

```
"%JAVA_HOME%\bin\javac" -classpath "%WAS_CLASSPATH%";  
<list of your application jars and classes> " -extdirs %WAS_EXT_DIRS%  
<your application class>.java
```

On UNIX systems, enter:

```
$JAVA_HOME/bin/javac -classpath "$WAS_CLASSPATH:"  
<list of your application jars and classes> " -extdirs $WAS_EXT_DIRS  
<your application class>.java
```

3. Run the following Java command to invoke your client application: On Windows systems, enter:

```
"%JAVA_HOME%\bin\java" "-Xbootclasspath/p:%WAS_BOOTCLASSPATH%"  
-Djava.security.auth.login.config="%WAS_HOME%\properties\wsjaas_client.conf"  
-classpath "%WAS_CLASSPATH%";<list of your application jars and classes>"  
-Djava.ext.dirs=%WAS_EXT_DIRS% -Djava.naming.provider.url=<an iiop URL or  
a corbaloc URL to your WebSphere server machine name>  
-Djava.naming.factory.initial=com.ibm.websphere.naming.WsnInitialContextFactory  
"%SERVER_ROOT%" "%CLIENTSAS%" <fully qualified class name to run>  
<your application parameters>
```

On UNIX systems, enter:

```
$JAVA_HOME/bin/java -Xbootclasspath/p:$WAS_BOOTCLASSPATH
-classpath "$WAS_CLASSPATH:<list of your application jars and classes>"
-Djava.ext.dirs=$WAS_EXT_DIRS -Djava.naming.provider.url=
<an IIOP URL or a Corbaloc URL to your WebSphere server machine name>
-Djava.naming.factory.initial=com.ibm.websphere.naming.WsnInitialContextFactory
"$SERVER_ROOT" "$CLIENTSAS" <fully qualified class name to run>
<your application parameters>
```

For more information on IIOP and corbaloc URLs, see Developing applications that use JNDI.

View the Samples gallery for more information about application clients. (See "Samples Gallery" in the Information Center.)

These samples do not include client applications that communicate with Enterprise server-side examples.

Developing thin application client code on a server machine

You must install WebSphere Application Server before performing this task.

1. Set the Java application thin client environment by using the **setupCmdLine** shell, located in:

```
install_root\bin\setupCmdLine.bat (on Windows systems)
install_root/bin/setupCmdLine.sh (on UNIX platforms)
```

2. Run the following Java compilation command to compile your client application: On Windows systems, enter:

```
"%JAVA_HOME%\bin\javac" -classpath "%WAS_CLASSPATH%";
<list of your application jars and classes> " -extdirs %WAS_EXT_DIRS%
<your application class>.java
```

On UNIX systems, enter:

```
$JAVA_HOME/bin/javac -classpath "$WAS_CLASSPATH:
<list of your application jars and classes>" -extdirs $WAS_EXT_DIRS
<your application class>.java
```

3. Run the application client. Perform one of the following methods:

- Use Java code to call your main class directly:

On Windows systems, enter:

```
"%JAVA_HOME%\bin\java" "-Xbootclasspath/p:%WAS_BOOTCLASSPATH%"
-Djava.security.auth.login.config="%WAS_HOME%\properties\wsjaas_client.conf"
-Djava.ext.dirs="%JAVA_HOME%\jre\lib\ext;%WAS_EXT_DIRS%"
-Djava.naming.provider.url=<an IIOP URL or a corbaloc URL to your
WebSphere server machine name>
-Djava.naming.factory.initial=com.ibm.websphere.naming.WsnInitialContextFactory
-Dserver.root="%WAS_HOME%" "%CLIENTSAS%" %USER_INSTALL_PROP%
-classpath "%WAS_CLASSPATH%;<list of your application jars and classes>"
<fully qualified class name to run><your application parameters>
```

On UNIX systems, enter:

```
"$JAVA_HOME/bin/java" "-Xbootclasspath/p:$WAS_BOOTCLASSPATH"
-Djava.security.auth.login.config="$WAS_HOME/properties/wsjaas_client.conf"
-Djava.ext.dirs="$JAVA_HOME/jre/lib/ext;%WAS_EXT_DIRS%"
-Djava.naming.provider.url=<an IIOP URL or a corbaloc URL to your
WebSphere server machine name>
-Djava.naming.factory.initial=com.ibm.websphere.naming.WsnInitialContextFactory
-Dserver.root="$WAS_HOME" $USER_INSTALL_PROP "$CLIENTSAS"
-classpath "$WAS_CLASSPATH;<list of your application jars and classes>
<fully qualified class name to run><your application parameters>
```

- Use the WebSphere Application Server launcher:

On Windows systems, enter:

```

"%JAVA_HOME%\bin\java" "-Xbootclasspath/p:%WAS_BOOTCLASSPATH%"
-Djava.security.auth.login.config="%WAS_HOME%\properties\wsjaas_client.conf"
"-Dws.ext.dirs=<list of your application jars and classes;
%WAS_EXT_DIRS%;%WAS_USER_DIRS%">
-Djava.naming.provider.url=<an IIOP URL or a corbaloc URL to your
WebSphere server machine name>
-Djava.naming.factory.initial=com.ibm.websphere.naming.WsnInitialContextFactory
"-Dserver.root=%WAS_HOME%"
"%CLIENTSAS%" %USER_INSTALL_PROP% -classpath "%WAS_CLASSPATH%"
com.ibm.ws.bootstrap.WSLauncher
<fully qualified class name to run><your application parameters>

```

On UNIX systems, enter:

```

"$JAVA_HOME/bin/java" "-Xbootclasspath/p:$WAS_BOOTCLASSPATH"
-Djava.security.auth.login.config="$WAS_HOME/properties/wsjaas_client.conf"
"-Dws.ext.dirs=<list of your application jars and classes>
$WAS_EXT_DIRS;$WAS_USER_DIRS"
-Djava.naming.provider.url=<an IIOP URL or a corbaloc URL to your
WebSphere server machine name>
-Djava.naming.factory.initial=com.ibm.websphere.naming.WsnInitialContextFactory
"-Dserver.root=$WAS_HOME"
"$CLIENTSAS" $USER_INSTALL_PROP -classpath "$WAS_CLASSPATH"
com.ibm.ws.bootstrap.WSLauncher
<fully qualified class name to run><your application parameters>

```

For more information on IIOP and corbaloc URLs, see *Developing applications that use JNDI*.

View the Samples gallery for more information about application clients. (See "Samples Gallery" in the Information Center.)

These samples do not include client applications that communicate with Enterprise server-side examples.

Assembling application clients

Assemble a client module to contain application client code. Group enterprise beans, Web components, and resource adapter code in separate modules.

Application client projects contain programs that run on networked client systems. An application client project is deployed as a JAR file.

Use the Assembly Toolkit to assemble an application client module in any of the following ways:

- Import an existing application client JAR file.
 - Create a new application client module.
1. Start the Assembly Toolkit.
 2. Optional: Open the J2EE perspective to work with J2EE projects. Click **Window > Open Perspective > Other > J2EE**.
 3. Optional: Open the J2EE Hierarchy view. Click **Window > Show View > J2EE Hierarchy**. Other helpful views include the Project Navigator view (**Window > Show View > Other > J2EE > Project Navigator**) and the Navigator view (**Window > Show View > Navigator**).
 4. **5.1+** Optional: To use application assembly services of WebSphere Business Integration (WBI) Server Foundation such as Application profiling, select the **Integration Server** target server.
 - a. Click **Project > Properties > J2EE**.
 - b. For **Target server**, select **Integration Server**. For this release, **Integration Server v5.1**.
 - c. Click **OK**.

You can set values for WBI Server Foundation services on the **Extended Services** tab of a deployment descriptor editor open on the module.

5. Migrate application client JAR files created with the Application Assembly Tool (AAT) or a different tool to the Assembly Toolkit. To migrate files, import your application client JAR files to the Assembly Toolkit.
6. Create a new application client.
7. Verify the contents of the new application client in either of the following ways:
 - In the J2EE Hierarchy view, expand **Application Client Modules** and view the new module.
 - Click **Window > Show View > Navigator** to see the associated files for the application client module in a Navigator view.

After you finish assembling all of your application's modules, you are ready to deploy your application.

Deploying application clients on workstation platforms

After developing an application client, deploy this application on client machines. *Deployment* consists of pulling together the various artifacts that the application client requires.

The *Application Client Resource Configuration Tool* (ACRCT) defines resources for the application client. These configurations are stored in the application client .ear file. The application client run time uses these configurations for resolving and creating an instance of the resources for the application client.

Note: This task only applies to J2EE application clients. Only perform this task if you configured your J2EE application client to use resource references.

1. Start the ACRCT and open an EAR file.
2. Configure new data source providers.
3. Configure mail providers and sessions.
4. Configure URL providers and sessions.
5. Configure Java messaging client resources.
6. Configure new environment entries.
7. (Optional) Remove application client resources.
8. Save the EAR file.

Starting the Application Client Resource Configuration Tool and opening an EAR file

Note: This task only applies to J2EE application clients.

1. Open a command prompt and change to the `install_root\bin` directory.
2. Run the `clientConfig.bat` file for a Windows system or the `clientConfig.sh` file for a UNIX system.
3. Open an EAR file within the Application Client Resource Configuration Tool (ACRCT):
 - Click **File > Open**.
 - Select the file and click **Open**.
4. Save your changes to the file and close the tool:
 - Click **File > Save**.
 - Click **File > Exit**.

Data sources for application clients

The J2EE application client does not support looking up or directly accessing data sources configured on WebSphere Application Server because the J2EE application client does not support Java 2 Connection Factories. To use a data source directly from the client application, you must use the ACRCT to configure your data source. In addition, WebSphere Application Server and WebSphere Application Server clients do not provide client database drivers to be used directly from a J2EE application client. If your application

client accesses a database directly, you must provide the database drivers on the client machine. You might contact your database vendor to acquire client database driver code and licenses. Instead of accessing the database directly, it is recommended that your client application use an enterprise bean. Accessing a database through an enterprise bean eliminates the need to have database drivers on the client machine, since the database access is handled by the enterprise bean running on the WebSphere Application Server. For a current list of providers that are supported on the WebSphere Application Server go the following site:

Supported hardware, software, and APIs

Configuring new data source providers (JDBC providers) for application clients

During this task, you create new data source providers, also known as JDBC providers, for your application client. In a separate administrative task, install the Java code for the required data source provider on the client machine on which the application client resides.

1. Start the tool and open the EAR file for which you want to configure the new data source provider. The EAR file contents display in a tree view.
2. Select the JAR file in which you want to configure the new data source provider from the tree.
3. Expand the JAR file to view its contents.
4. Click the **Data Source Providers** folder. Do one of the following:
 - Right-click the folder and click **New Provider**.
 - Click **Edit > New** on the menu bar.
5. Configure the data source provider properties in the resulting property dialog.
6. Click **OK** when you finish.
7. Click **File > Save** on the menu bar to save your changes.

Configuring new data source providers

During this task, you will create new data source providers, also known as JDBC drivers, for your application client. In a separate administrative task, install the Java code for the required data source provider on the client machine where the application client resides.

1. Start the Application Client Resource Configuration Tool (ACRCT).
2. Click **File > Open**, and select the EAR file for which you want to configure the new data source provider. The EAR file contents display in a tree view.
3. Select the JAR file in which you want to configure the new data source provider from the tree.
4. Expand the JAR file to view its contents.
5. Right click the **Data Source Providers** folder and select **New Provider**.
6. Configure the data source provider properties in the resulting property dialog.
7. Click **OK**.
8. Click **File > Save** to save your changes.

Example: Configuring data source provider and data source settings: The purpose of this article is to help you to configure data source provider and data source settings.

- Required fields:
 - Data Source Provider Properties page: name
 - Data Source Properties page: name, jndiName
- Special cases:
 - The user name and password fields have no equivalent XML tags. You must specify these fields in the custom properties.
 - The password is encrypted when you use the Application Client Resource Configuration Tool (ACRCT). If you do not use the ACRCT the field cannot be encrypted.

- **Example:**

```
<resources.jdbc:JDBCProvider xmi:id="JDBCProvider_1" name="jdbcProvider:name"
description="jdbcProvider:description" implementationClassName="jdbcProvider:
ImplementationClass">
<classpath>jdbcProvider:classpath</classpath>
<factories xmi:type="resources.jdbc:WAS40DataSource" xmi:id="WAS40DataSource_1"
name="jdbcFactory:name" jndiName="jdbcFactory:jndiName"
description="jdbcFactory:description" databaseName="jdbcFactory:databasename">
<propertySet xmi:id="J2EEResourcePropertySet_13">
<resourceProperties xmi:id="J2EEResourceProperty_13" name="jdbcFactory:customName"
value="jdbcFactory:customValue"/>
<resourceProperties xmi:id="J2EEResourceProperty_14" name="user"
value="jdbcFactory:user"/>
<resourceProperties xmi:id="J2EEResourceProperty_15" name="password"
value="{xor}NTs9PBk+PCswLSZ1MT4y0g==" />
</propertySet>
</factories>
<propertySet xmi:id="J2EEResourcePropertySet_14">
<resourceProperties xmi:id="J2EEResourceProperty_16" name="jdbcProvider:customName"
value="jdbcProvider:customeValue"/>
</propertySet>
</resources.jdbc:JDBCProvider>
```

Data source provider settings for application clients:

Use this page to create a data source under a JDBC provider which provides the specific JDBC driver implementation class.

To view this Application Client Resource Configuration Tool (ACRCT) page, click **File > Open**. After you browse for an EAR file, click **Open**. Expand the selected JAR file. Right-click **Data Source Providers >** and click **New**. The following fields appear on the **General** tab:

Name:

Specifies the display name for the data source.

For example you can set this field to *Test Data Source*.

Data type String

Description:

Specifies a text description for the resource.

Data type String

Class Path:

A list of paths or .jar file names which together form the location for the resource provider classes.

Implementation class:

Use this setting to perform database specific functions.

Data type String
Default Dependent on JDBC driver implementation class

Custom Properties:

Specifies name-value pairs for setting additional properties on the object that is created at run time for this resource.

You must enter a name that is a public property on the object and a value that can be converted from a string to the type required by the set method of the property. The acceptable properties and values depend on the object that is created. Refer to the object documentation for a list of valid properties and values.

Data source properties for application clients:

Use this page to create or modify the data sources.

To view this Application Client Resource Configuration Tool (ACRCT) page, click **File > Open**. After you browse for an EAR file, click **Open**. Expand the selected JAR file > **Data Source Providers > Data source provider instance**. Right-click **Data Sources** and click **New**. The following fields are displayed on the **General** tab:

Name:

Specifies the display name of this data source.

Data type String

Description:

Specifies a text description of the data source.

Data type String

JNDI Name:

The application client run time uses this field to retrieve configuration information.

Database Name:

The name of the database to which you want to connect.

User:

Use the user ID with the Password property, for authentication if the calling application does not provide a user ID and password explicitly.

If you specify a value for the User ID property, you must also specify a value for the Password property. The connection factory User ID and Password properties are used if the calling application does not provide a user ID and password explicitly.

Password:

Use the password with the User ID property, for authentication if the calling application does not provide a user ID and password explicitly.

If you specify a value for the User ID property, you must also specify a value for the Password property.

Re-Enter Password:

Confirms the password.

Custom Properties:

Specifies name-value pairs for setting additional properties on the object that is created at run time for this resource.

You must enter a name that is a public property on the object and a value that can be converted from a string to the type required by the set method of the property. The acceptable properties and values depend on the object that is created. Refer to the object documentation for a list of valid properties and values.

Configuring new data sources for application clients

During this task, you create new data sources for your application client.

1. Click the data source provider for which you want to create a data source in the tree. Take one of the following actions as needed:
 - Configure a new data source provider.
 - Click an existing data source provider.
2. Expand the data source provider to view its **Data Sources** folder.
3. Click the data source folder. Take one of the following actions as needed:
 - Right click the data source folder and click **New Factory**.
 - Click **Edit > New** on the menu bar.
4. Configure the data source properties in the displayed fields.
5. Click **OK** when you finish.
6. Click **File > Save** on the menu bar to save your changes.

Configuring mail providers and sessions for application clients

Use the Application Client Resource Configuration Tool (ACRCT) to edit the configurations of JavaMail sessions and providers for your application clients to use.

1. Open the ACRCT.
2. Open an EAR file.
3. Locate the JavaMail objects in the tree that displays. For example, if your file contains JavaMail sessions, expand **Resources > application.jar > JavaMail Providers > java_mail_provider_instance > JavaMail Sessions**.

In this example, *java_mail_provider_instance* is a particular JavaMail provider.

The JavaMail session instances are located in the **JavaMail Sessions** folder.

Mail provider settings for application clients

Use this page to implement the JavaMail API and create mail sessions.

To view this Application Client Resource Configuration Tool (ACRCT) page, click **File > Open**. After you browse for an EAR file, click **Open**. Expand the selected JAR file. Right-click **Mail Providers >** and click **New**. The following fields appear on the **General** tab:

Name:

The name of the JavaMail resource provider.

Description:

An optional description for the resource provider.

Class Path:

Specifies a list of paths or JAR file names which together form the location for the resource provider classes.

Protocol:

Specifies the name of the protocol.

Classname:

Specifies the name of the class implementing the protocol. Leave this field blank if you want to use the default implementation.

Type:

This menu contains the following two values: TRANSPORT or STORE.

Custom Properties:

Specifies name-value pairs for setting additional properties on the object that is created at run time for this resource.

You must enter a name that is a public property on the object and a value that can be converted from a string to the type required by the set method of the property. The acceptable properties and values depend on the object that is created. Refer to the object documentation for a list of valid properties and values.

Mail session settings for application clients

Use this page to configure mail session properties.

To view this Application Client Resource Configuration Tool (ACRCT) page, click **File > Open**. After you browse for an EAR file, click **Open**. Expand the selected JAR file > **Mail Providers > mail provider instance**. Right-click **Mail Sessions** and click **New**. The following fields appear on the **General** tab:

Name:

Represents the administrative name of the JavaMail session object.

Description:

Provides an optional description for your administrative records.

JNDI Name:

The application client run time uses this field to retrieve configuration information. The name must match the value of the **Name** field on the General tab in the Application Client Resource Reference section of the Assembly Tool.

Mail Transport Host:

Specifies the server to connect to when sending mail.

Mail Transport Protocol:

Specifies the transport protocol to use when sending mail.

Mail Transport User:

Specifies the user ID to use when the mail transport host requires authentication.

Mail Transport Password:

Specifies the password to use when the mail transport host requires authentication.

Re-Enter Password:

Confirms the password.

Mail From:

Specifies the mail originator.

Mail Store Host:

Specifies the mail account host (or "domain") name.

Mail Store User:

Specifies the user ID of the mail account.

Mail Store Password:

Specifies the password of the mail account.

Re-Enter Password:

Confirms the password.

Mail Store Protocol:

Specifies the protocol to be used when receiving mail.

Mail Debug:

When true, JavaMail interaction with mail servers, along with these mail session properties are printed to the `stdout` file.

Custom Properties:

Specifies name-value pairs for setting additional properties on the object that is created at run time for this resource.

You must enter a name that is a public property on the object and a value that can be converted from a string to the type required by the set method of the property. The acceptable properties and values depend on the object that is created. Refer to the object documentation for a list of valid properties and values.

Example: Configuring JavaMail provider and JavaMail session settings for application clients

The purpose of this article is to help you configure JavaMail provider and JavaMail session settings.

- Required fields:
 - JavaMail Provider Properties page: name, and at least one protocol provider
 - JavaMail Session Properties page: name, jndiName, mail transport protocol, mail store protocol
- Special cases:
 - The password is encrypted when using the ACRCT tool. Without the tool, you cannot encrypt this field.
- Example:

```

<resources.mail:MailProvider xmi:id="MailProvider_1" name="Default Mail Provider"
description="IBM JavaMail Implementation">
<classpath>mailProvider:classpath</classpath>
<factories xmi:type="resources.mail:MailSession" xmi:id="MailSession_1"
name="mailSession:name" jndiName="mailSession:jndiName"
description="mailSession:description" mailTransportHost="mailSession:mailTransportHost"
mailTransportUser="mailSession:mailTransportUser"
mailTransportPassword="{xor}Mj42Mww6LCw2MDF1MT4y0g=="
mailFrom="mailSession:mailFrom" mailStoreHost="mailSession:mailStoreHost"
mailStoreUser="mailSession:mailStoreUser"
mailStorePassword="{xor}Mj42Mww6LCw2MDF1MT4y0g==" debug="true"
mailTransportProtocol="ProtocolProvider_1" mailStoreProvider="ProtocolProvider_1">
<propertySet xmi:id="J2EEResourcePropertySet_1">
<resourceProperties xmi:id="J2EEResourceProperty_1"
name="mailSession:customName" value="mailSession:customValue"/>
</propertySet>
</factories>
<propertySet xmi:id="J2EEResourcePropertySet_2">
<resourceProperties xmi:id="J2EEResourceProperty_2" name="mailProvider:customName"
value="mailProvider:customValue"/>
</propertySet>
<protocolProviders xmi:id="ProtocolProvider_1" protocol="smtp"
classname="smtp:className"/>
<protocolProviders xmi:id="ProtocolProvider_2" protocol="pop3"
classname="pop3:className"/>
<protocolProviders xmi:id="ProtocolProvider_3" protocol="imap"
classname="imap:className"/>
</resources.mail:MailProvider>

```

Configuring new mail sessions for application clients

During this task, you configure new mail sessions for your application client. The mail sessions are associated with the pre-configured default mail provider supplied by the product.

1. Start the Application Client Resource Configuration Tool (ACRCT) and open the EAR file. The EAR file contents are displayed in a tree view.
2. Select the JAR file in which you want to configure the new JavaMail session.
3. Expand the JAR file to view its contents.
4. Click **JavaMail Providers** > **MailProvider** > **JavaMail Sessions**. Complete one of the following actions:
 - Right click the **JavaMail Sessions** folder and select **New Factory**.
 - Click **Edit** > **New** on the menu bar.
5. Configure the JavaMail session properties in the displayed fields.
6. Click **OK**.
7. Click **File** > **Save** on the menu bar to save your changes.

URLs for application clients

A *Uniform Resource Locator* (URL) is an identifier that points to an electronically accessible resource, such as a directory file on a machine in a network, or a document stored in a database.

URLs appear in the format *scheme:scheme_information*.

You can represent a *scheme* as http, ftp, file, or another term that identifies the type of resource and the mechanism by which you can access the resource.

In a World Wide Web browser location or address box, a URL for a file available using HyperText Transfer Protocol (HTTP) starts with http:. An example is http://www.ibm.com. Files available using File Transfer Protocol (FTP) start with ftp:. Files available locally start with file:.

The *scheme_information* commonly identifies the Internet machine making a resource available, the path to that resource, and the resource name. The *scheme_information* for HTTP, FTP and File generally starts with two slashes (//), then provides the Internet address separated from the resource path name with one slash (/). For example,

```
http://www-4.ibm.com/software/webservers/appserv/library.html.
```

For HTTP and FTP, the path name ends in a slash when the URL points to a directory. In such cases, the server generally returns the default index for the directory.

URL providers for the Application Client Resource Configuration Tool

A URL provider implements the function for a particular URL protocol, such as Hyper Text Transfer Protocol (HTTP). This provider, comprised of a pair of classes, extends the `java.net.URLStreamHandler` and `java.net.URLConnection` classes.

Configuring new URL providers for application clients

During this task, you create URL providers and URLs for your client application. In a separate administrative task, you must install the Java code for the required URL provider on the client machine on which the client application resides.

1. Start the Application Client Resource Configuration Tool (ACRCT).
2. Open the EAR file for which you want to configure the new URL provider. The EAR file contents display in a tree view.
3. Select the JAR file in which you want to configure the new URL provider from the tree.
4. Expand the JAR file to view the contents.
5. Click the folder called **URL Providers**. Complete one of the following actions:
 - Right click the folder and select **New Provider**.
 - Click **Edit > New** on the menu bar.
6. Configure the URL provider properties in the resulting property dialog.
7. Click **OK**.
8. Click **File > Save** on the menu bar to save your changes.

Configuring URL providers and sessions using the Application Client Resource Configuration Tool

Use the Application Client Resource Configuration Tool (ACRCT) to edit the configurations of URL providers and URLs to be used by your application clients.

1. Open the ACRCT.
2. Open an EAR file.
3. Locate the URL objects in the tree that displays. For example, if your file contains URL providers and URLs, expand **Resources** -> **application.jar** -> **URL Providers** -> **url_provider_instance** where **url_provider_instance** is a particular URL provider.
4. If you expand the tree further, you will also see the **URLs** folders containing the URL instances for each URL provider instance.

URL settings for application clients:

Use this page to implement the function for a particular URL protocol, such as Hyper Text Transfer Protocol (HTTP).

To view this Application Client Resource Configuration Tool (ACRCT) page, click **File > Open**. After you browse for an EAR file, click **Open**. Expand the selected JAR file > **URL Providers > URL provider instance**. Right-click **URLs** and click **New**. The following fields appear on the **General** tab.

This provider, comprised of classes, extends the `java.net.URLStreamHandler` and `java.net.URLConnection` classes.

Name:

Administrative name for the URL

Description:

Optional description of the URL, for your administrative records

JNDI Name:

The application client run time uses this field to retrieve configuration information. The name must match the value of the **Name** field on the General tab in the Application Client Resource Reference section of the Application Assembly Tool.

URL:

A Uniform Resource Locator (URL) name that points to an Internet or intranet resource. For example:
`http://www.ibm.com`.

Custom Properties:

Specifies name-value pairs for setting additional properties on the object that is created at run time for this resource.

You must enter a name that is a public property on the object and a value that can be converted from a string to the type required by the set method of the property. The acceptable properties and values depend on the object that is created. Refer to the object documentation for a list of valid properties and values.

URL provider settings for application clients:

Use this page create new URL providers.

To view this Application Client Resource Configuration Tool (ACRCT) page, click **File > Open**. After you browse for an EAR file, click **Open**. Expand the selected JAR file. Right click **URL Providers**, and click **New**. The following fields appear on the **General** tab.

A URL provider implements the function for a particular URL protocol, such as Hyper Text Transfer Protocol (HTTP). This provider, comprised of classes, extends the `java.net.URLStreamHandler` and `java.net.URLConnection` classes.

Name:

Administrative name for the URL

Description:

Optional description of the URL, for your administrative records

Class Path:

A list of paths or JAR file names which together form the location for the resource provider classes.

Protocol:

Protocol supported by this stream handler. For example, "nntp", "smtp", "ftp", and so on.

To use the default protocol, leave this field blank.

Stream handler class:

Fully qualified name of a User-defined Java class that extends the `java.net.URLStreamHandler` for a particular URL protocol, such as FTP.

To use the default stream handler, leave this field blank.

Custom Properties:

Specifies name-value pairs for setting additional properties on the object that is created at run time for this resource.

You must enter a name that is a public property on the object and a value that can be converted from a string to the type required by the set method of the property. The acceptable properties and values depend on the object that is created. Refer to the object documentation for a list of valid properties and values.

Example: Configuring URL and URL provider settings for application clients

The purpose of this article is to help you to configure URL and URL provider settings.

- Required fields:
 - URL Properties page: name, jndiName, url
 - URL Provider Properties page: name
- Example:

```
<resources.url:URLProvider xmi:id="URLProvider_1" name="urlProvider:name"
description="urlProvider:description"
streamHandlerClassName="urlProvider:streamHandlerClass"
protocol="urlProvider:protocol">
<classpath>urlProvider:classpath</classpath>
<factories xmi:type="resources.url:URL" xmi:id="URL_1" name="urlFactory:name"
jndiName="urlFactory:jndiName" description="urlFactory:description"
spec="urlFactory:url">
<propertySet xmi:id="J2EEResourcePropertySet_18">
<resourceProperties xmi:id="J2EEResourceProperty_20" name="urlFactory:customName"
value="urlFactory:customValue"/>
</propertySet>
</factories>
<propertySet xmi:id="J2EEResourcePropertySet_19">
<resourceProperties xmi:id="J2EEResourceProperty_21" name="urlProvider:customName"
value="urlProvider:customValue"/>
</propertySet>
</resources.url:URLProvider>
```

Configuring new URLs with the Application Client Resource Configuration Tool

During this task, you create URLs for your client application.

1. Click the URL provider for which you want to create a URL in the tree. Do one of the following:
 - Configure a new URL provider.
 - Click an existing URL provider.
2. Expand the URL provider to view the **URLs** folder.
3. Click the URL folder. Complete one of the following actions:

- Right click the folder and click **New Factory**.
 - Click **Edit** -> **New** on the menu bar.
4. Configure the URL properties in the displayed fields.
 5. Click **OK** when you finish.
 6. Click **File** > **Save** in the menu bar to save your changes.

WebSphere asynchronous messaging using the Java Message Service API for the Application Client Resource Configuration Tool

WebSphere Application Server supports asynchronous messaging as a method of communication based on the Java Message Service (JMS) programming interface. The JMS interface provides a common way for Java programs (clients and J2EE applications) to create, send, receive, and read asynchronous requests as JMS messages.

This topic provides an overview of asynchronous messaging using JMS support provided by the WebSphere Application Server.

The base support for asynchronous messaging using the JMS API provides the common set of JMS interfaces and associated semantics that define how a JMS client can access the facilities of a JMS provider. This support enables WebSphere product J2EE applications, as JMS clients, to exchange messages asynchronously with other JMS clients, by using JMS destinations (queues or topics). A J2EE application can use JMS queue destinations for point-to-point messaging and JMS topic destinations for Publisher and Subscriber messaging. A J2EE application can explicitly poll for messages on a destination, and then retrieve messages for processing by business logic beans (enterprise beans).

With the base JMS and XA support, the J2EE application uses standard JMS calls to process messages, including any responses or outbound messaging. An enterprise bean can handle responses acting as a sender bean, or within the enterprise bean that receives the incoming messages. Optionally, this process can use two-phase commit within the scope of a transaction. This level of function for asynchronous messaging is called *bean-managed messaging*, and gives an enterprise bean complete control over the messaging infrastructure, for example, connection and session pool management. The common container has no role in bean-managed messaging.

WebSphere Application Server also supports automatic asynchronous messaging using message-driven beans (a type of enterprise bean defined in the EJB 2.0 specification) and JMS listeners (part of the JMS application server facilities). Messages are automatically retrieved from JMS destinations, optionally within a transaction, then sent to the message-driven bean in a J2EE application, without the application having to explicitly poll JMS destinations.

Configuring Java messaging client resources

In a separate administrative task, install the Java Message Service (JMS) client on the client machine where the application client resides. The messaging product vendor must provide an implementation of the JMS client. For more information, see your messaging product documentation.

During this task, you create new JMS provider configurations for your application client. The application client can use a messaging service through the Java Message Service APIs. A JMS provider provides two kinds of J2EE factories. One is a *JMS connection factory*, and the other is a *JMS destination factory*.

1. Start the Application Client Resource Configuration Tool (ACRCT).
2. Open the EAR file for which you want to configure the new JMS provider. The EAR file contents are in the displayed tree view.
3. Select the JAR file in which you want to configure the new JMS provider from the tree.
4. Expand the JAR file to view its contents.
5. Click the **JMS Providers** folder and click **New Provider**.

6. Configure the JMS provider properties in the resulting property dialog.
7. Click **OK**.
8. Click **File > Save**.

Configuring new JMS providers with the Application Client Resource Configuration Tool

During this task, you will create new JMS provider configurations for your application client. The application client makes use of a messaging service through the Java Message Service (JMS) APIs. A JMS provider provides two kinds of J2EE factories. One is a JMS Connection factory, and the other is a JMS destination factory.

In a separate administrative task, you must install the JMS client on the client machine where the application client resides. The messaging product vendor must provide an implementation of the JMS client. For more information, see your messaging product documentation.

1. Start the tool and open the EAR file for which you want to configure the new JMS provider. The EAR file contents will be displayed in a tree view.
2. From the tree, select the JAR file in which you want to configure the new JMS provider.
3. Expand the JAR file to view its contents.
4. Click the **JMS Providers** folder. Complete one of the following actions:
 - Right click the folder and select **New Provider**.
 - On the menu bar, click **Edit > New**.
5. In the resulting property dialog, configure the JMS provider properties.
6. Click **OK** when finished.
7. Click **File -> Save** on the menu bar to save your changes.

JMS provider settings for application clients

Use this page to configure properties of the Java Messaging Service (JMS) provider, if you want to use a JMS provider other than the internal WebSphere product JMS provider or the MQSeries product JMS provider.

To view this Application Client Resource Configuration Tool (ACRCT) page, click **File > Open**. After you browse for an EAR file, click **Open**. Expand the selected JAR file. Right click **JMS Providers**, and click **New**. The following fields appear on the **General** tab.

Name:

The name by which the JMS provider is known for administrative purposes.

Data type String

Description:

A description of the JMS provider, for administrative purposes

Data type String

Class Path:

A list of paths or .jar file names which together form the location for the resource provider classes.

Context factory class:

The Java class name of the initial context factory for the JMS provider.

For example, for an LDAP service provider the value has the form: `com.sun.jndi.ldap.LdapCtxFactory`.

Data type String

Provider URL:

The JMS provider URL for external JNDI lookups.

For example, an LDAP URL for a JMS provider has the form: `ldap://hostname.company.com/contextName`.

Data type String

Custom Properties:

Specifies name-value pairs for setting additional properties on the object that is created at run time for this resource.

You must enter a name that is a public property on the object and a value that can be converted from a string to the type required by the set method of the property. The acceptable properties and values depend on the object that is created. Refer to the object documentation for a list of valid properties and values.

Version 5 Default Provider queue connection factory settings for application clients

Use this panel to view or change the configuration properties of the selected queue connection factory for use with the internal WebSphere Application Server product Java Messaging Service (JMS) provider that is installed with WebSphere Application Server. These configuration properties control how connections are created to the associated JMS queue destination.

To view this Application Client Resource Configuration Tool (ACRCT) page, click **File > Open**. After you browse for an EAR file, click **Open**. Expand the selected JAR file > **JMS Providers > JMS provider instance**. Right-click **WAS Queue Connection Factories** and click **New**. The following fields appear on the **General** tab.

A queue connection factory is used to create JMS connections to queue destinations. The queue connection factory is created by the internal WebSphere Application Server product JMS provider. A queue connection factory for the internal WebSphere Application Server product JMS provider has the following properties:

Name:

The name by which this queue connection factory is known for administrative purposes within IBM WebSphere Application Server. The name must be unique within the JMS connection factories across the WebSphere administrative domain.

Data type String

Description:

A description of this connection factory for administrative purposes within IBM WebSphere Application Server.

Data type String
Default Null

JNDI Name:

The application client run time uses this field to retrieve configuration information. The name must match the value of the Name field on the General tab in the Application Client Resource Reference section of the Application Assembly Tool.

User:

The User ID used, with the **Password** property, for authentication if the calling application does not provide a `userid` and password explicitly.

If you specify a value for the **User ID** property, you must also specify a value for the **Password** property.

The connection factory **User ID** and **Password** properties are used if the calling application does not provide a User ID and password explicitly, for example, if the calling application uses the method `createQueueConnection()`. The JMS client flows the `userid` and password to the JMS server.

Data type String

Password:

The password used, with the **User ID** property, for authentication if the calling application does not provide a `userid` and password explicitly.

If you specify a value for the **User ID** property, you must also specify a value for the **Password** property.

Re-Enter Password:

Confirms the password.

Note:

The WebSphere node name of the administrative node where the JMS server runs for this connection factory. Connections created by this factory connect to that JMS server.

Data type String

Application Server:

Enter the name of the application server. This name is not the host name of the machine, but the name of the configured application server.

Custom Properties:

Specifies name-value pairs for setting additional properties on the object that is created at run time for this resource.

You must enter a name that is a public property on the object and a value that can be converted from a string to the type required by the set method of the property. The acceptable properties and values depend on the object that is created. Refer to the object documentation for a list of valid properties and values.

Version 5 Default Provider topic connection factory settings for application clients

Use this panel to view or change the configuration properties of the selected topic connection factory for use with the internal WebSphere Application Server product Java Messaging Service (JMS) provider. These configuration properties control how connections are created to the associated JMS topic destination.

To view this Application Client Resource Configuration Tool (ACRCT) page, click **File > Open**. After you browse for an EAR file, click **Open**. Expand the selected JAR file > **JMS Providers > JMS provider instance**. Right click **WAS Topic Connection Factories** and click **New**. The following fields appear on the **General** tab.

A topic connection factory is used to create JMS connections to topic destinations. The topic connection factory is created by the associated JMS provider. A topic connection factory for the internal WebSphere product JMS provider has the following properties.

Name:

The name by which this queue connection factory is known for administrative purposes within IBM WebSphere Application Server. The name must be unique within the JMS connection factories across the WebSphere administrative domain.

Data type String

Description:

A description of this topic connection factory for administrative purposes within IBM WebSphere Application Server.

Data type String
Default Null

JNDI Name:

The application client run time uses this field to retrieve configuration information. The name must match the value of the **Name** field on the General tab in the Application Client Resource Reference section of the Assembly Tool.

User:

The user ID used, with the **Password** property, for authentication if the calling application does not provide a userid and password explicitly.

If you specify a value for the **User ID** property, you must also specify a value for the **Password** property.

The connection factory **User ID** and **Password** properties are used if the calling application does not provide a userid and password explicitly, for example, if the calling application uses the method `createTopicConnection()`. The JMS client flows the userid and password to the JMS server.

Data type String

Password:

The password used, with the **User ID** property, for authentication if the calling application does not provide a userid and password explicitly.

If you specify a value for the **User ID** property, you must also specify a value for the **Password** property.

Data type	String
Default	Null

Re-Enter Password:

Confirms the password.

Node:

The WebSphere node name of the administrative node where the JMS server runs for this connection factory. Connections created by this factory connect to that JMS server.

Data type	Enum
Default	Null
Range	Pull-down list of nodes in the WebSphere administrative domain.

Application Server:

Enter the name of the application server. This name is not the host name of the machine, but the name of the configured application server.

Port:

Which of the two ports that connections use to connect to the JMS Server. The QUEUED port is for full-function JMS publish/subscribe support, the DIRECT port is for non-persistent, non-transactional, non-durable subscriptions only.

Note: Message-driven beans cannot use the direct listener port for publish/subscribe support. Therefore, any topic connection factory configured with the Port set to Direct cannot be used with message-driven beans.

Data type	Enum
Units	Not applicable
Default	QUEUED
Range	QUEUED The listener port used for full-function JMS-compliant, publish/subscribe support. DIRECT The listener port used for direct TCP/IP connection (non-transactional, non-persistent, and non-durable subscriptions only) for publish/subscribe support. The TCP/IP port numbers for these ports are defined on the product internal JMS server.

Client ID:

The JMS client identifier used for connections to the MQSeries queue manager.

Data type	String
------------------	--------

Custom Properties:

Specifies name-value pairs for setting additional properties on the object that is created at run time for this resource.

You must enter a name that is a public property on the object and a value that can be converted from a string to the type required by the set method of the property. The acceptable properties and values depend on the object that is created. Refer to the object documentation for a list of valid properties and values.

Version 5 Default Provider queue destination settings for application clients

Use this panel to view or change the configuration properties of the selected queue destination for use with WebSphere Application Server product Java Messaging Service (JMS) provider.

To view this Application Client Resource Configuration Tool (ACRCT) page, click **File > Open**. After you browse for an EAR file, click **Open**. Expand the selected JAR file > **JMS Providers > JMS provider instance**. Right click **WAS Queue Destinations** and click **New**. The following fields are displayed on the **General** tab.

A queue destination is used to configure the properties of a JMS queue. Connections to the queue are created by the associated queue connection factory for the internal product JMS provider. A queue for use with the internal product JMS provider has the following properties.

Name:

The name by which the queue is known for administrative purposes within IBM WebSphere Application Server.

Data type	String
------------------	--------

Description:

A description of the queue, for administrative purposes

Data type	String
Default	Null

JNDI Name:

The application client run time uses this field to retrieve configuration information. The name must match the value of the **Name** field on the General tab in the Application Client Resource Reference section of the Application Assembly Tool.

Persistence:

Whether all messages sent to the destination are persistent, non-persistent, or have their persistence defined by the application

Data type	Enum
Units	Not applicable
Default	APPLICATION_DEFINED

Range

Application defined

Messages on the destination have their persistence defined by the application that put them onto the queue.

Queue defined

[WebSphere MQ destination only] Messages on the destination have their persistence defined by the WebSphere MQ queue definition properties.

Persistent

Messages on the destination are persistent.

Non persistent

Messages on the destination are not persistent.

Priority:

Whether the message priority for this destination is defined by the application or the **Specified priority** property

Data type
Units
Default
Range

Enum
Not applicable
APPLICATION_DEFINED
Application defined
The priority of messages on this destination is defined by the application that put them onto the destination.
Queue defined
[WebSphere MQ destination only] Messages on the destination have their persistence defined by the WebSphere MQ queue definition properties.
Specified
The priority of messages on this destination is defined by the **Specified priority** property. *If you select this option, you must define a priority on the **Specified priority** property.*

Specified Priority:

If the **Priority** property is set to Specified, type here the message priority for this queue, in the range 0 (lowest) through 9 (highest)

If the **Priority** property is set to Specified, messages sent to this queue have the priority value specified by this property.

Data type
Units
Default
Range

Integer
Message priority level
Null
0 (lowest priority) through 9 (highest priority)

Expiry:

Whether the expiry timeout for this queue is defined by the application or the **Specified expiry** property, or whether messages on the queue expire (have an unlimited expiry timeout).

Data type
Units
Default

Enum
Not applicable
APPLICATION_DEFINED

Range**Application defined**

The expiry timeout for messages in this queue is defined by the application that put them onto the queue.

Specified

The expiry timeout for messages in this queue is defined by the **Specified expiry** property. If you select this option, you must define a time out on the **Specified expiry** property.

Unlimited

Messages in this queue have no expiry timeout, and those messages never expire.

Specified Expiry:

If the **Expiry timeout** property is set to *Specified*, specify the number of milliseconds (greater than 0) after which messages on this queue expire.

Data type

Integer

Units

Milliseconds

Default

Null

Range

Greater than or equal to 0

- 0 indicates that messages never timeout
- Other values are an integer number of milliseconds

Custom Properties:

Specifies name-value pairs for setting additional properties on the object that is created at run time for this resource.

You must enter a name that is a public property on the object and a value that can be converted from a string to the type required by the set method of the property. The acceptable properties and values depend on the object that is created. Refer to the object documentation for a list of valid properties and values.

Version 5 Default Provider topic destination settings for application clients

Use this panel to view or change the configuration properties of the selected topic destination for use with the internal WebSphere Application Server product Java Messaging Service (JMS) provider.

To view this Application Client Resource Configuration Tool (ACRCT) page, click **File > Open**. After you browse for an EAR file, click **Open**. Expand the selected JAR file > **JMS Providers > JMS provider instance**. Right click **WAS Topic Destinations** and click **New**. The following fields appear on the **General** tab.

A topic destination is used to configure the properties of a JMS topic for the associated JMS provider. Connections to the topic are created by the associated topic connection factory. A topic used with the internal JMS provider has the following properties.

Name:

The name by which the topic is known for administrative purposes.

Data type

String

Description:

A description of the topic, for administrative purposes within IBM WebSphere Application Server.

Data type	String
Default	Null

JNDI Name:

The application client run-time environment uses this field to retrieve configuration information. The name must match the value of the **Name** field on the General tab in the Application Client Resource Reference section of the Application Assembly Tool.

Topic Name: The name of the topic as defined to the JMS provider.

Data type	String
------------------	--------

Persistence:

Whether all messages sent to the destination are persistent, non-persistent, or have their persistence defined by the application

Data type	Enum
Units	Not applicable
Default	APPLICATION_DEFINED
Range	Application defined Messages on the destination have their persistence defined by the application that put them onto the queue. Queue defined [WebSphere MQ destination only] Messages on the destination have their persistence defined by the WebSphere MQ queue definition properties. Persistent Messages on the destination are persistent. Non persistent Messages on the destination are not persistent.

Priority:

Whether the message priority for this destination is defined by the application or the **Specified priority** property.

Data type	Enum
Units	Not applicable
Default	APPLICATION_DEFINED

Range**Application defined**

The priority of messages on this destination is defined by the application that put them onto the destination.

Queue defined

[WebSphere MQ destination only] Messages on the destination have their persistence defined by the WebSphere MQ queue definition properties.

Specified

The priority of messages on this destination is defined by the **Specified priority** property. *If you select this option, you must define a priority on the **Specified priority** property.*

Specified Priority:

If the **Priority** property is set to Specified, specify the message priority for this queue, in the range 0 (lowest) through 9 (highest).

If the **Priority** property is set to Specified, messages sent to this queue have the priority value specified by this property.

Data type	Integer
Units	Message priority level
Default	Null
Range	0 (lowest priority) through 9 (highest priority)

Expiry:

Whether the expiry timeout for this queue is defined by the application or the **Specified expiry** property, or messages on the queue never expire (have an unlimited expiry timeout)

Data type	Enum
Units	Not applicable
Default	APPLICATION_DEFINED
Range	Application defined The expiry timeout for messages on this queue is defined by the application that put them onto the queue. Specified The expiry timeout for messages on this queue is defined by the Specified expiry property. <i>If you select this option, you must define a timeout on the Specified expiry property.</i> Unlimited Messages on this queue have no expiry timeout, so those messages never expire.

Specified Expiry:

If the **Expiry timeout** property is set to Specified, type here the number of milliseconds (greater than 0) after which messages on this queue expire

Data type	Integer
Units	Milliseconds
Default	Null

Range

Greater than or equal to 0

- 0 indicates that messages never time out
- Other values are an integer number of milliseconds

Custom Properties:

Specifies name-value pairs for setting additional properties on the object that is created at run time for this resource.

You must enter a name that is a public property on the object and a value that can be converted from a string to the type required by the set method of the property. The acceptable properties and values depend on the object that is created. Refer to the object documentation for a list of valid properties and values.

WebSphere MQ Provider queue connection factory settings for application clients

Use this panel to view or change the configuration properties of the selected queue connection factory for use with the MQSeries product Java Messaging Service (JMS) provider. These configuration properties control how connections are created to the associated JMS queue destination.

To view this Application Client Resource Configuration Tool (ACRCT) page, click **File > Open**. After you browse for an EAR file, click **Open**. Expand the selected JAR file > **JMS Providers > JMS provider instance**. Right click **MQ Queue Connection Factories**, and click **New**. The following fields are displayed on the **General** tab.

A queue connection factory creates JMS connections to queue destinations. The queue connection factory is created by the MQSeries product JMS provider. A queue connection factory for the JMS provider has the following properties.

Note:

- The property values that you specify must match the values that you specified when configuring MQSeries for JMS resources. For more information about configuring MQSeries product JMS resources, see the MQSeries *Using Java* book, located in the WebSphere MQ Family library.
- In MQSeries, names can have a maximum of 48 characters, with the exception of channels which have a maximum of 20 characters.

Name:

The name by which this queue connection factory is known for administrative purposes within IBM WebSphere Application Server. The name must be unique within the JMS connection factories across the WebSphere administrative domain.

Data type

String

Description:

A description of this connection factory for administrative purposes within IBM WebSphere Application Server.

Data type

String

Default

Null

JNDI Name:

The application client run time uses this field to retrieve configuration information. The name must match the value of the **Name** field on the General tab in the Application Client Resource Reference section of the Assembly Tool.

User:

The user ID used, with the **Password** property, for authentication if the calling application does not provide a userid and password explicitly.

If you specify a value for the **User ID** property, you must also specify a value for the **Password** property.

The connection factory **User ID** and **Password** properties are used if the calling application does not provide a userid and password explicitly; for example, if the calling application uses the method `createQueueConnection()`. The JMS client flows the userid and password to the JMS server.

Data type String

Password:

The password used, with the **User ID** property, for authentication if the calling application does not provide a userid and password explicitly.

If you specify a value for the **User ID** property, you must also specify a value for the **Password** property.

Data type String
Default Null

Re-Enter Password:

Confirms the password.

Queue Manager:

The name of the MQSeries queue manager for this connection factory.

Connections created by this factory connect to that queue manager.

Data type String

Host:

The name of the host on which the WebSphere MQ queue manager runs for client connection only.

Data type String
Default Null
Range A valid TCP/IP host name

Port:

The TCP/IP port number used for connection to the WebSphere MQ queue manager, for client connection only.

This port must be configured on the WebSphere MQ queue manager.

Data type	Integer
Default	Null
Range	A valid TCP/IP port number, configured on the WebSphere MQ queue manager.

Channel:

The name of the channel used for connection to the WebSphere MQ queue manager, for client connection only.

Data type	String
Default	Null
Range	1 through 20 ASCII characters

Transport type:

Specifies whether the WebSphere MQ client connection or JNI bindings are used for connection to the WebSphere MQ queue manager. The external JMS provider controls the communication protocols between JMS clients and JMS servers. Tune the transport type when you are using non-ASF nonpersistent, nondurable, nontransactional messaging or when you want to satisfy security issues and the client is local to the queue manager node.

Data type	Enum
Units	Not applicable
Default	BINDINGS
Range	BINDINGS JNI bindings are used to connect to the queue manager. BINDINGS is a shared memory protocol and can only be used when the queue manager is on the same node as the JMS client and poses security risks that should be addressed through the use of EJB roles. CLIENT WebSphere MQ client connection is used to connect to the queue manager. CLIENT is a typical TCP-based protocol. DIRECT For WebSphere MQ Event Broker using DIRECT mode. DIRECT is a lightweight sockets protocol used in nontransactional, nondurable and nonpersistent Publish/Subscribe messaging. DIRECT is only works for clients and message-driven beans using the non-ASF protocol. QUEUED QUEUED is a standard TCP protocol.

Recommended

Queue connection factory transport type

BINDINGS is faster by 30% or more, but it lacks security. When you have security concerns, BINDINGS is more desirable than CLIENT.

Topic connection factory transport type

DIRECT is the fastest type and should be used where possible. Use BINDINGS when you want to satisfy additional security tasks and the queue manager is local to the JMS client. QUEUED is the fallback for all other cases. WebSphere MQ 5.3 before CSD2 with the DIRECT setting can lose messages when used with message-driven beans and under load. This loss also happens with client-side applications unless the broker maxClientQueueSize is set to 0. You can set this to 0 with the command (shown on multiple line for publication):

```
#wempschangeproperties WAS_nodeName_server1 -e default -o
DynamicSubscriptionEngine -n maxClientQueueSize -v 0 -x executionGroupUUID
```

where executionGroupUUID can be found by starting the broker and looking in the Event Log/Applications for event 2201. This value is usually ffffffff-0000-0000-000000000000.

Client ID:

The JMS client identifier used for connections to the MQSeries queue manager.

Data type String

CCSID:

The coded character set identifier for use with the WebSphere MQ queue manager.

This coded character set identifier (CCSID) must be one of the CCSIDs supported by WebSphere MQ.

Data type String

For more information about supported CCSIDs, and about converting between message data from one coded character set to another, see the *WebSphere MQ System Administration* and the *WebSphere MQ Application Programming Reference* books. These references are available from the WebSphere MQ messaging multiplatform and platform-specific books Web pages; for example, at <http://www-3.ibm.com/software/ts/mqseries/library/manualsa/manuals/platspecific.html>, the IBM Publications Center, or from the WebSphere MQ collection kit, SK2T-0730.

Message Retention:

Select this check box to specify that unwanted messages are to be left on the queue. Otherwise, unwanted messages are handled according to their disposition options.

Data type	Enum
Units	Not applicable
Default	Cleared
Range	Selected Unwanted messages are left on the queue.
	Cleared Unwanted messages are handled according to their disposition options.

Custom Properties:

Specifies name-value pairs for setting additional properties on the object that is created at run time for this resource.

You must enter a name that is a public property on the object and a value that can be converted from a string to the type required by the set method of the property. The acceptable properties and values depend on the object that is created. Refer to the object documentation for a list of valid properties and values.

WebSphere MQ Provider topic connection factory settings for application clients

Use this panel to view or change the configuration properties of the selected topic connection factory for use with the MQSeries product Java Messaging Service (JMS) provider. These configuration properties control how connections are created to the associated JMS topic destination.

To view this Application Client Resource Configuration Tool (ACRCT) page, click **File > Open**. After you browse for an EAR file, click **Open**. Expand the selected JAR file > **JMS Providers > JMS provider instance**. Right-click **MQ Topic Connection Factories** and click **New**.

A topic connection factory is used to create JMS connections to topic destinations. The topic connection factory is created by the MQSeries product JMS provider. A topic connection factory for the MQSeries product JMS provider has the following properties.

Note:

- The property values that you specify must match the values that you specified when configuring MQSeries product JMS resources. For more information about configuring MQSeries product JMS resources, see the MQSeries *Using Java* book.
- In MQSeries, names can have a maximum of 48 characters, with the exception of channels which have a maximum of 20 characters.

Name:

The name by which this topic connection factory is known for administrative purposes within IBM WebSphere Application Server. The name must be unique within the JMS provider.

Data type String

Description:

A description of this topic connection factory for administrative purposes within IBM WebSphere Application Server.

Data type String

Default Null

JNDI Name:

The Java Naming and Directory Interface (JNDI) name that is used to bind the topic connection factory into the application server name space.

As a convention, use the fully qualified JNDI name; for example, in the form *jms/Name*, where *Name* is the logical name of the resource.

This name is used to link the platform binding information. The binding associates the resources defined by the deployment descriptor of the module to the actual (physical) resources bound into JNDI by the platform.

Data type String

Units	En_US ASCII characters
Default	Null
Range	1 through 45 ASCII characters

User:

The user ID used, with the **Password** property, for authentication if the calling application does not provide a userid and password explicitly.

If you specify a value for the **User** property, you must also specify a value for the **Password** property.

The connection factory **User** and **Password** properties are used if the calling application does not provide a userid and password explicitly, for example, if the calling application uses the method createTopicConnection(). The JMS client flows the userid and password to the JMS server.

Data type	String
------------------	--------

Password:

The password used, with the **User ID** property, for authentication if the calling application does not provide a userid and password explicitly.

If you specify a value for the **User ID** property, you must also specify a value for the **Password** property.

Data type	String
Default	Null

Re-Enter Password:

Confirms the password.

Queue Manager:

The name of the MQSeries queue manager for this connection factory. Connections created by this factory connect to that queue manager.

Data type	String
------------------	--------

Host:

The name of the host on which the WebSphere MQ queue manager runs for client connections only.

Data type	String
Default	Null
Range	A valid TCP/IP host name

Port:

The TCP/IP port number used for connection to the WebSphere MQ queue manager, for client connection only.

This port must be configured on the WebSphere MQ queue manager.

Data type	Integer
Default	Null
Range	A valid TCP/IP port number, configured on the WebSphere MQ queue manager.

Channel:

The name of the channel used for client connections to the WebSphere MQ queue manager for client connection only.

Data type	String
Default	Null
Range	1 through 20 ASCII characters

Transport Type:

Whether MQSeries client connection or JNDI bindings are used for connection to the MQSeries queue manager.

Data type	Enum
Units	Not applicable
Default	BINDINGS
Range	CLIENT MQSeries client connection is used to connect to the MQSeries queue manager. BINDINGS JNDI bindings are used to connect to the MQSeries queue manager.

Client ID:

The JMS client identifier used for connections to the MQSeries queue manager.

Data type	String
------------------	--------

CCSID:

The coded character set identifier to be used with the WebSphere MQ queue manager.

This coded character set identifier (CCSID) must be one of the CCSIDs supported by WebSphere MQ.

Data type	String
------------------	--------

Broker Control Queue:

The name of the broker control queue to which all command messages (except publications and requests to delete publications) are sent.

The name of the broker control queue. Publisher and subscriber applications as well as other brokers send all command messages (except publications and requests to delete publications) to this queue.

Data type	String
------------------	--------

Units	En_US ASCII characters
Default	Null
Range	1 through 48 ASCII characters

Broker Queue Manager:

The name of the MQSeries queue manager that provides the Publisher and Subscriber message broker.

Data type	String
Units	En_US ASCII characters
Default	Null
Range	1 through 48 ASCII characters

Broker Pub Queue:

The name of the broker input queue that receives all publication messages for the default stream

The name of the broker's input queue (stream queue) that receives all publication messages for the default stream. Applications can also send requests to delete publications on the default stream to this queue.

Data type	String
Units	En_US ASCII characters
Default	Null
Range	1 through 48 ASCII characters

Broker Sub Queue:

The name of the broker queue from which nondurable subscription messages are retrieved.

The name of the broker queue from which nondurable subscription messages are retrieved. The subscriber specifies the name of the queue when it registers a subscription.

Data type	String
Units	En_US ASCII characters
Default	Null
Range	1 through 48 ASCII characters

Broker CCSubQ:

The name of the broker queue from which nondurable subscription messages are retrieved for a ConnectionConsumer request. This property applies only for use of the Web container.

Data type	String
Units	En_US ASCII characters
Default	Null
Range	1 through 48 ASCII characters

Broker Version:

Specifies whether the message broker is provided by the MQSeries MA0C SupportPac or newer versions of WebSphere family message broker products.

Data type	Enum
Units	Not applicable
Default	Advanced
Range	<p>Advanced The message broker is provided by newer versions of WebSphere family message broker products (MQ Integrator and MQ Publish and Subscribe)</p> <p>Basic The message broker is provided by the MQSeries MA0C SupportPac (MQSeries - Publish and Subscribe)</p>

Custom Properties:

Specifies name-value pairs for setting additional properties on the object that is created at run time for this resource.

You must enter a name that is a public property on the object and a value that can be converted from a string to the type required by the set method of the property. The acceptable properties and values depend on the object that is created. Refer to the object documentation for a list of valid properties and values.

WebSphere MQ Provider queue destination settings for application clients

Use this panel to view or change the configuration properties of the selected queue destination for use with the MQSeries product Java Messaging Service (JMS) provider.

To view this Application Client Resource Configuration Tool (ACRCT) page, click **File > Open**. After you browse for an EAR file, click **Open**. Expand the selected JAR file > **JMS Providers > JMS provider instance**. Right-click **MQ Queue Destinations** and click **New**. The following fields are displayed on the **General** tab.

A queue destination configures the properties of a JMS queue. Connections to the queue are created by the associated queue connection factory for the MQSeries product JMS provider. A queue for use with the MQSeries product JMS provider has the following properties.

Note:

- The property values that you specify must match the values that you specified when configuring MQSeries product JMS resources. For more information about configuring MQSeries product JMS resources, see the MQSeries *Using Java* book.
- In MQSeries, names can have a maximum of 48 characters.

Name:

The name by which the queue is known for administrative purposes within IBM WebSphere Application Server.

Data type	String
------------------	--------

Description:

A description of the queue, for administrative purposes

Data type	String
Default	Null

JNDI Name:

The application client run-time environment uses this field to retrieve configuration information. The name must match the value of the **Name** field on the General tab in the Application Client Resource Reference section of the Application Assembly Tool.

Persistence:

Whether all messages sent to the destination are persistent, nonpersistent or have their persistence defined by the application.

Data type	Enum
Units	Not applicable
Default	APPLICATION_DEFINED
Range	<p>Application defined Messages on the destination have their persistence defined by the application that put them onto the queue.</p> <p>Queue defined [WebSphere MQ destination only] Messages on the destination have their persistence defined by the WebSphere MQ queue definition properties.</p> <p>Persistent Messages on the destination are persistent.</p> <p>Nonpersistent Messages on the destination are not persistent.</p>

Priority:

Whether the message priority for this destination is defined by the application or the **Specified priority** property

Data type	Enum
Units	Not applicable
Default	APPLICATION_DEFINED
Range	<p>Application defined The priority of messages on this destination is defined by the application that put them onto the destination.</p> <p>Queue defined [WebSphere MQ destination only] Messages on the destination have their persistence defined by the WebSphere MQ queue definition properties.</p> <p>Specified The priority of messages on this destination is defined by the Specified priority property. <i>If you select this option, you must define a priority on the Specified priority property.</i></p>

Specified Priority:

If the **Priority** property is set to Specified, specify the message priority for this queue, in the range 0 (lowest) through 9 (highest).

Data type	Integer
Units	Message priority level
Default	Null
Range	0 (lowest priority) through 9 (highest priority)

Expiry:

Whether the expiry timeout value for this queue is defined by the application or the by **Specified expiry** property or whether messages on the queue never expire (have an unlimited expiry timeout).

Data type	Enum
Units	Not applicable
Default	APPLICATION_DEFINED
Range	Application defined The expiry timeout for messages on this queue is defined by the application that put them onto the queue. Specified The expiry timeout for messages on this queue is defined by the Specified expiry property. If you select this option, you must define a timeout on the Specified expiry property. Unlimited Messages on this queue have no expiry timeout and those messages never expire.

Specified Expiry:

If the **Expiry timeout** property is set to **Specified**, type here the number of milliseconds (greater than 0) after which messages on this queue expire.

Data type	Integer
Units	Milliseconds
Default	Null
Range	Greater than or equal to 0 <ul style="list-style-type: none">• 0 indicates that messages never time out• Other values are an integer number of milliseconds

Base Queue Name:

The name of the queue to which messages are sent, on the queue manager specified by the **Base queue manager name** property

Data type	String
------------------	--------

Base Queue Manager Name:

The name of the MQSeries queue manager to which messages are sent

This queue manager provides the queue specified by the **Base queue name** property.

Data type	String
Units	En_US ASCII characters
Default	Null
Range	A valid MQSeries Queue Manager name, as 1 through 48 ASCII characters

CCSID:

The coded character set identifier to use with the WebSphere MQ queue manager.

This coded character set identifier (CCSID) must be one of the CCSID identifier supported by WebSphere MQ queue manager.

Data type String

Integer encoding:

If native encoding is not enabled, select whether integer encoding is normal or reversed.

Data type Enum
Units Not applicable
Default NORMAL
Range **NORMAL**
Normal integer encoding is used.
REVERSED
Reversed integer encoding is used.

For more information about encoding properties, see the WebSphere MQ *Using Java* document.

Decimal encoding:

Indicates that if native encoding is not enabled to select whether decimal encoding is normal or reversed.

Data type Enum
Units Not applicable
Default NORMAL
Range **NORMAL**
Normal decimal encoding is used.
REVERSED
Reversed decimal encoding is used.

For more information about encoding properties, see the WebSphere MQ *Using Java* document.

Floating point encoding:

Indicates that if native encoding is not enabled to select the type of floating point encoding.

Data type Enum
Units Not applicable
Default IEEEENORMAL
Range **IEEEENORMAL**
IEEE normal floating point encoding is used.
IEEEEVERSED
IEEE reversed floating point encoding is used.
S390 S390 floating point encoding is used.

For more information about encoding properties, see the WebSphere MQ *Using Java* document.

Native encoding:

Indicates that the queue destination use native encoding (appropriate encoding values for the Java platform) when you select this check box.

Data type
Units
Default
Range

Enum
Not applicable
Cleared
Cleared

Native encoding is not used, so specify the following properties for integer, decimal and floating point encoding.

Selected

Native encoding is used (to provide appropriate encoding values for the Java platform).

For more information about encoding properties, see the MQSeries *Using Java* document.

Target client:

Whether the receiving application is JMS-compliant or is a traditional WebSphere MQ application

Data type
Units
Default
Range

Enum
Not applicable
MQSeries
MQSeries

The target is a traditional WebSphere MQ application that does not support JMS.

JMS

The target application supports JMS.

Custom Properties:

Specifies name-value pairs for setting additional properties on the object that is created at run time for this resource.

You must enter a name that is a public property on the object and a value that can be converted from a string to the type required by the set method of the property. The acceptable properties and values depend on the object that is created. Refer to the object documentation for a list of valid properties and values.

WebSphere MQ Provider topic destination settings for application clients

Use this panel to view or change the configuration properties of the selected topic destination for use with the MQSeries product Java Messaging Service (JMS) provider.

To view this Application Client Resource Configuration Tool (ACRCT) page, click **File > Open**. After you browse for an EAR file, click **Open**. Expand the selected JAR file > **JMS Providers > JMS provider instance**. Right click **MQ Topic Destinations**, and click **New**. The following fields are displayed on the **General** tab.

A topic destination is used to configure the properties of a JMS topic for the associated JMS provider. Connections to the topic are created by the associated topic connection factory. A topic for use with the MQSeries product JMS provider has the following properties.

Note:

- The property values that you specify must match the values that you specified when configuring MQSeries product JMS resources. For more information about configuring MQSeries product JMS resources, see the MQSeries *Using Java* book.
- In MQSeries, names can have a maximum of 48 characters.

Name:

The name by which the topic is known for administrative purposes.

Data type String

Description:

A description of the topic for administrative purposes within IBM WebSphere Application Server.

JNDI Name:

The application client run time uses this field to retrieve configuration information. The name must match the value of the **Name** field on the General tab in the Application Client Resource Reference section of the Application Assembly Tool.

Persistence:

Specifies whether all messages sent to the destination are persistent, nonpersistent, or have their persistence defined by the application.

Data type	Enum
Units	Not applicable
Default	APPLICATION_DEFINED
Range	Application defined Messages on the destination have their persistence defined by the application that put them in the queue. Queue defined [WebSphere MQ destination only] Messages on the destination have their persistence defined by the WebSphere MQ queue definition properties. Persistent Messages on the destination are persistent. Nonpersistent Messages on the destination are not persistent.

Priority:

Specifies whether the message priority for this destination is defined by the application or the **Specified priority** property.

Data type	Enum
Units	Not applicable
Default	APPLICATION_DEFINED
Range	Application defined The priority of messages on this destination is defined by the application that put them in the destination. Queue defined [WebSphere MQ destination only] Messages on the destination have their persistence defined by the WebSphere MQ queue definition properties. Specified The priority of messages on this destination is defined by the Specified priority property. If you select this option, you must define a priority for the Specified priority property.

Specified Priority:

If the **Priority** property is set to *Specified*, type the message priority for this queue, in the range 0 (lowest) through 9 (highest).

If the **Priority** property is set to *Specified*, messages sent to this queue have the priority value specified by this property.

Data type	Integer
Units	Message priority level
Default	Null
Range	0 (lowest priority) through 9 (highest priority)

Expiry:

Whether the expiry timeout for this queue is defined by the application or by the **Specified expiry** property or by messages on the queue never expire (have an unlimited expiry timeout).

Data type	Enum
Units	Not applicable
Default	APPLICATION_DEFINED
Range	Application defined The expiry timeout for messages on this queue is defined by the application that put them in the queue. Specified The expiry timeout for messages in this queue is defined by the Specified expiry property. If you select this option, you must define a timeout value for the Specified expiry property. Unlimited Messages on this queue have no expiry timeout, and these messages never expire.

Specified Expiry:

If the **Expiry timeout** property is set to *Specified*, type the number of milliseconds (greater than 0) after which messages on this queue expire.

Data type	Integer
Units	Milliseconds
Default	Null
Range	Greater than or equal to 0 <ul style="list-style-type: none">• 0 indicates that messages never time out• Other values are an integer number of milliseconds

Base Topic Name:

The name of the topic to which messages are sent.

Data type	String
------------------	--------

CCSID:

The coded character set identifier to use with the WebSphere MQ queue manager.

This coded character set identifier (CCSID) must be one of the CCSID identifiers that WebSphere MQ supports.

Data type String

Integer encoding:

Indicates whether integer encoding is normal or reversed when native encoding is not enabled.

Data type	Enum
Units	Not applicable
Default	NORMAL
Range	NORMAL Normal integer encoding is used. REVERSED Reversed integer encoding is used.

For more information about encoding properties, see the WebSphere MQ *Using Java* document.

Decimal encoding:

If native encoding is not enabled, select whether decimal encoding is normal or reversed.

Data type	Enum
Units	Not applicable
Default	NORMAL
Range	NORMAL Normal decimal encoding is used. REVERSED Reversed decimal encoding is used.

For more information about encoding properties, see the WebSphere MQ *Using Java* document.

Floating point encoding:

Indicates the type of floating point encoding when native encoding is not enabled.

Data type	Enum
Units	Not applicable
Default	IEEEENORMAL
Range	IEEEENORMAL IEEE normal floating point encoding is used. IEEEEVERSED IEEE reversed floating point encoding is used. S390 S/390 floating point encoding is used.

For more information about encoding properties, see the WebSphere MQ *Using Java* document.

Native encoding:

Indicates that the queue destination uses native encoding (appropriate encoding values for the Java platform) when you select this check box.

Data type	Enum
Units	Not applicable
Default	Cleared
Range	<p>Cleared Native encoding is not used, so specify the previous properties for integer, decimal and floating point encoding.</p> <p>Selected Native encoding is used (to provide appropriate encoding values for the Java platform).</p> <p>For more information about encoding properties, see the MQSeries <i>Using Java</i> document.</p>

BrokerDurSubQueue:

The name of the broker queue from which durable subscription messages are retrieved.

The subscriber specifies the name of the queue when it registers a subscription.

Data type	String
Units	En_US ASCII characters
Default	Null
Range	1 through 48 ASCII characters

BrokerCCDurSubQueue:

The name of the broker queue from which durable subscription messages are retrieved for a ConnectionConsumer. This property applies only for use of the Web container.

Data type	String
Units	En_US ASCII characters
Default	Null
Range	1 through 48 ASCII characters

Target Client:

Specifies whether the receiving application is JMS compliant or is a traditional MQSeries application.

Data type	Enum
Units	Not applicable
Default	MQSeries
Range	<p>MQSeries The target is a traditional MQSeries application that does not support JMS.</p> <p>JMS The target is a JMS compliant application.</p>

Custom Properties:

Specifies the name-value pairs for setting additional properties on the object that is created at run time for this resource.

You must enter a name that is a public property on the object and a value that can be converted from a string to the type required by the set method of the property. The acceptable properties and values depend on the object that is created. Refer to the object documentation for a list of valid properties and values.

Generic JMS connection factory settings for application clients

Use this panel to view or change the configuration properties of the selected JMS connection factory for use with the associated JMS provider. These configuration properties control how connections are created to the associated JMS destination.

To view this Application Client Resource Configuration Tool (ACRCT) page, click **File > Open**. After you browse for an EAR file, click **Open**. Expand the selected JAR file > **JMS Providers > new_JMS_Provider_instance**. Right click **JMS Connection Factories**, and click **New**. The following fields are displayed on the **General** tab.

A Java Message Service (JMS) connection factory creates connections to JMS destinations. The JMS connection factory is created by the associated JMS provider. A JMS connection factory for a generic JMS provider (other than the internal WebSphere product JMS provider or the MQSeries product JMS provider) has the following properties:

Name:

The name by which this JMS connection factory is known for administrative purposes within IBM WebSphere Application Server. The name must be unique within the associated JMS provider.

Data type String

Description:

A description of this connection factory for administrative purposes within IBM WebSphere Application Server.

Data type String

Default Null

JNDI Name:

The application client run time uses this field to retrieve configuration information. The name must match the value of the **Name** field on the General tab in the Application Client Resource Reference section of the Application Assembly Tool.

User:

Indicates the user ID used with the **Password** property, for authentication if the calling application does not provide a userid and password explicitly.

If you specify a value for the **User ID** property, you must also specify a value for the **Password** property.

The connection factory **User ID** and **Password** properties are used if the calling application does not provide a userid and password explicitly; for example, if the calling application uses the method `createQueueConnection()`. The JMS client flows the userid and password to the JMS server.

Data type String

Password:

The password used with the **User ID** property for authentication if the calling application does not provide a userid and password explicitly.

If you specify a value for the **User ID** property, you must also specify a value for the **Password** property.

Data type	String
Default	Null

Re-Enter Password:

Confirms the password entered in the **Password** field.

External JNDI Name:

The JNDI name that is used to bind the queue into the application server name space.

As a convention, use the fully qualified JNDI name, for example, *jms/Name*, where *Name* is the logical name of the resource.

This name is used to link the platform binding information. The binding associates the resources defined by the deployment descriptor of the module to the actual (physical) resources bound into JNDI API by the platform.

Data type	String
------------------	--------

Connection Type:

Whether this JMS destination is a queue (for point-to-point) or topic (for pub/sub).

Select one of the following options:

Queue

A JMS queue destination for point-to-point messaging.

Topic A JMS topic destination for publish subscribe messaging.

Custom Properties:

Specifies name-value pairs for setting additional properties on the object that is created at run time for this resource.

You must enter a name that is a public property on the object and a value that can be converted from a string to the type required by the set method of the property. The acceptable properties and values depend on the object that is created. Refer to the object documentation for a list of valid properties and values.

Generic JMS destination settings for application clients

Use this panel to view or change the configuration properties of the selected JMS destination for use with the associated JMS provider.

To view this Application Client Resource Configuration Tool (ACRCT) page, click **File > Open**. After you browse for an EAR file, click **Open**. Expand the selected JAR file > **JMS Providers > new JMS Provider instance**. Right-click **JMS Destinations**, and click **New**. The following fields are displayed on the **General** tab.

A JMS destination is used to configure the properties of a JMS destination for the associated generic JMS provider. Connections to the JMS destination are created by the associated JMS connection factory. A JMS destination for use with a generic JMS provider (not the internal WebSphere product JMS provider or MQSeries product JMS provider) has the following properties.

Name:

The name by which the queue is known for administrative purposes within IBM WebSphere Application Server.

Data type String

Description:

A description of the queue, for administrative purposes

JNDI Name:

The JNDI name of the actual (physical) name of the JMS destination bound into JNDI.

External JNDI Name:

The JNDI name that is used to bind the queue into the application server name space.

As a convention, use the fully qualified JNDI name; for example, in the form *jms/Name*, where *Name* is the logical name of the resource.

This name is used to link the platform binding information. The binding associates the resources defined by the deployment descriptor of the module to the actual (physical) resources bound into JNDI by the platform.

Data type String

Destination Type:

Whether this JMS destination is a queue (for point-to-point) or topic (for publishing and subscribing).

Select one of the following options:

Queue

A JMS queue destination for point-to-point messaging.

Topic A JMS topic destination for pub/sub messaging.

Custom Properties:

Specifies name-value pairs for setting additional properties on the object that is created at run time for this resource.

You must enter a name that is a public property on the object and a value that can be converted from a string to the type required by the set method of the property. The acceptable properties and values depend on the object that is created. Refer to the object documentation for a list of valid properties and values.

Example: Configuring JMS provider, JMS connection factory and JMS destination settings for application clients

The purpose of this article is to help you to configure JMS Provider, JMS Connection Factory and JMS Destination settings.

- Required fields:
 - JMS Provider Properties page: name, and at least one protocol provider
 - JMS Connection Factory Properties page: name, jndiName, destination type
 - JMS Destination Properties page: name, jndiName, destination type
- Special cases:
 - The destination type must be QUEUE, or TOPIC.
- Example:

```

<resources.jms:JMSProvider xmi:id="JMSProvider_3" name="genericJMSProvider:name"
description="genericJMSProvider:description"
externalInitialContextFactory="genericJMSProvider:contextFactoryClass"
externalProviderURL="genericJMSProvider:providerUrl">
<classpath>genericJMSProvider:classpath</classpath>
<factories xmi:type="resources.jms:GenericJMSDestination"
xmi:id="GenericJMSDestination_1" name="jmsDestination:name"
jndiName="jmsDestination:jndiName" description="jmsDestination:description"
externalJNDIName="jmsDestination:externalJndiName" type="QUEUE">
<propertySet xmi:id="J2EEResourcePropertySet_15">
<resourceProperties xmi:id="J2EEResourceProperty_17" name="jmsDestination:customName"
value="jmsDestination:customValue"/>
</propertySet>
</factories>
<factories xmi:type="resources.jms:GenericJMSConnectionFactory"
xmi:id="GenericJMSConnectionFactory_1" name="jmsCF:name" jndiName="jmsCF:jndiName"
description="jmsCF:description" userID="jmsCF:user" password="{xor}NTIsHB11MT4y0g=="
externalJNDIName="jmsCF:externalJndiName" type="QUEUE">
<propertySet xmi:id="J2EEResourcePropertySet_16">
<resourceProperties xmi:id="J2EEResourceProperty_18" name="jmsCF:customName"
value="jmsCF:customValue"/>
</propertySet>
</factories>
<propertySet xmi:id="J2EEResourcePropertySet_17">
<resourceProperties xmi:id="J2EEResourceProperty_19"
name="genericJMSProvider:customName" value="genericJMSProvider:customValue"/>
</propertySet>
</resources.jms:JMSProvider>

```

Configuring new connection factories for application clients

During this task, you create a new Java Messaging Service (JMS) connection factory configuration for your application client.

1. Click the JMS provider for which you want to create a connection factory in the tree. Complete one of the following actions:
 - Configure a new JMS provider.
 - Click an existing JMS provider.
2. Expand the JMS provider to view its **JMS Connection Factories** folder.
3. Click the connection factory folder, and complete one of the following actions:
 - Right-click the folder and select **New Factory**.
 - Click **Edit > New** on the menu bar.
4. Configure the JMS connection factory properties in the displayed fields.
5. Click **OK** when you finish.
6. Click **File > Save** on the menu bar to save your changes.

Configuring new Java Message Service destinations for application clients

During this task, you create a new Java Message Service (JMS) destination configuration for your application client.

1. Click the JMS provider in the tree for which you want to create a destination. Complete one of the following actions:
 - Configure a new JMS provider.
 - Click an existing JMS provider.
2. Expand the JMS provider to view its **JMS Destinations** folder.
3. Click the provider folder, and complete one of the following actions:
 - Right-click the folder and select **New Factory**.

- Click **Edit > New** on the menu bar.
- 4. Configure the JMS destination properties in the displayed fields.
- 5. Click **OK** when you finish.
- 6. Click **File > Save** on the menu bar to save your changes.

Example: Configuring MQ Queue and Topic connection factories and destination factories for application clients

The purpose of this article is to help you configure MQ Queue connection factory, MQ Topic connection factory, MQ Queue destination factory, and MQ Topic destination factory settings.

- Required fields:
 - MQ Queue Connection Factory Properties page: name, jndiName and transport type
 - MQ Topic Connection Factory Properties page: name, jndiName and broker Version
 - MQ Queue Factory Properties page: name, jndiName, persistence, priority, expiry, baseQueueName and targetClient
 - MQ Topic Factory Properties page: name, jndiName, persistence, priority, expiry, baseQueueName and targetClient
- Special cases:
 - The transport type must be CLIENT, or BINDINGS.
 - The Broker Version must be MA0C, or MQSI.
 - The port must be a numerical value between -2417483648 and 2417483647.
 - The CCSID must be a numerical value between -2417483648 and 2417483647.
 - The persistence value must be APPLICATION_DEFINED, QUEUE_DEFINED, PERSISTENT or, NONPERSISTENT.
 - The priority must be APPLICATION_DEFINED, QUEUE_DEFINED, or SPECIFIED.
 - The expiry must be APPLICATION_DEFINED, UNLIMITED, or SPECIFIED.
 - The integer encoding must be Normal, or Reversed.
 - The decimal encoding must be Normal, or Reversed.
 - The floating encoding must be IEEENormal, IEEEReversed or S390.
 - The target client must be JMS or MQ.
 - On the MQ Queue Connection Factory Properties page, only set the queueManager, host, and port values. These are required fields if the transport type is CLIENT.
 - On the MQ Topic Connection Factory Properties page, only set the queueManager, host, and port (required) fields if the transport type is CLIENT.
 - On the MQ Topic Factory Properties, and the MQ Queue Factory Properties pages, only set the Integer encoding, decimal encoding, and floating point encoding (required) fields if you do not set the nativeEncoding value.
 - On the MQ Topic Factory Properties and the MQ Queue Factory Properties pages, the specified priority entry field must be an integer between 0 and 9 if priority is set to SPECIFIED .
 - On the MQ Topic Factory Properties and the MQ Queue Factory Properties pages, the specified expiry entry field must be a value greater than 0 if the expiry value is set to SPECIFIED.
- Example:

```
<resources.jms:JMSProvider xmi:id="JMSProvider_1" name="MQ JMS Provider"
description="mqJMSProvider:description"
externalInitialContextFactory="mqJMSProvider:contextFactoryClass"
externalProviderURL="mqJMSProvider:providerUrl">
<classpath>mqJMSProvider:classpath</classpath>
<factories xmi:type="resources.jms.mqseries:MQQueueConnectionFactory"
xmi:id="MQQueueConnectionFactory_1" name="mqQCF:name" jndiName="mqQCF:jndiName"
description="mqQCF:description" userID="mqQCF:user" password="{xor}Mi40HB1lMT4y0g=="
queueManager="mqQCF:queueManager" host="mqQCF:host" port="1" channel="mqQCF:channel"
transportType="CLIENT" clientID="mqQCF:clientId" CCSID="2">
<propertySet xmi:id="J2EEResourcePropertySet_3">
<resourceProperties xmi:id="J2EEResourceProperty_3" name="mqQCF:customName"
value="mqQCF:customValue"/>
</propertySet>
</factories>
<factories xmi:type="resources.jms.mqseries:MQTopicConnectionFactory"
xmi:id="MQTopicConnectionFactory_1" name="mqTCF:name" jndiName="mqTCF:jndiName">
```

```

description="mqTCF:description" userID="mqTCF:user"
password="{xor}Mi4LHB1lNTE7NhE+Mjo=" host="mqTCF:host" port="1"
transportType="CLIENT" channel="mqTCF:channel" queueManager="mqTCF:queueManager"
brokerControlQueue="mqTCF:brokerControlQueue"
brokerQueueManager="mqTCF:brokerQueueManager" brokerPubQueue="mqTCF:brokerPubQueue"
brokerSubQueue="mqTCF:brokerSubQueue" brokerCCSubQ="mqTCF:brokerCCSubQ"
brokerVersion="MAQC" clientID="mqTCF:clientId" CCSID="2">
<propertySet xmi:id="J2EEResourcePropertySet_4">
<resourceProperties xmi:id="J2EEResourceProperty_4" name="mqTCF:customName"
value="mqTCF:customValue"/>
</propertySet>
</factories>
<factories xmi:type="resources.jms.mqseries:MQQueue" xmi:id="MQQueue_1" name="mqQ:name"
jndiName="mqQ:jndiName" description="mqQ:description" persistence="APPLICATION_DEFINED"
priority="SPECIFIED" specifiedPriority="1" expiry="SPECIFIED" specifiedExpiry="1"
baseQueueName="mqQ:baseQueueName" baseQueueManagerName="mqQ:baseQueueManagerName"
CCSID="1" integerEncoding="Normal" decimalEncoding="Normal"
floatingPointEncoding="IEEENormal" targetClient="JMS">
<propertySet xmi:id="J2EEResourcePropertySet_5">
<resourceProperties xmi:id="J2EEResourceProperty_5" name="mqQ:customName"
value="mqQ:customValue"/>
</propertySet>
</factories>
<factories xmi:type="resources.jms.mqseries:MQTopic" xmi:id="MQTopic_1"
name="mqT:name" jndiName="mqT:jndiName" description="mqT:description"
persistence="APPLICATION_DEFINED" priority="SPECIFIED" specifiedPriority="1"
expiry="SPECIFIED" specifiedExpiry="2" baseTopicName="mqT:baseTopicName" CCSID="3"
integerEncoding="Normal" decimalEncoding="Normal" floatingPointEncoding="IEEENormal"
targetClient="JMS" brokerDurSubQueue="mqT:brokerDurSubQueue"
brokerCCDurSubQueue="mqT:brokerCCDurSubQueue">
<propertySet xmi:id="J2EEResourcePropertySet_6">
<resourceProperties xmi:id="J2EEResourceProperty_6" name="mqT:customName"
value="mqT:customValue"/>
</propertySet>
</factories>
<propertySet xmi:id="J2EEResourcePropertySet_7">
<resourceProperties xmi:id="J2EEResourceProperty_7" name="mqJMSProvider:customName"
value="mqJMSProvider:customValue"/>
</propertySet>
</resources.jms:JMSProvider>

```

Example: Configuring WAS Queue and Topic connection factories and destination factories for application clients

The purpose of this article is to help you configure WAS Queue connection factory, WAS Topic connection factory, WAS Queue destination factory, and WAS Topic destination factory settings.

- Required fields:
 - JMS Provider Properties page: name
 - WAS Queue Connection Factory Properties page: name, jndiName and node
 - WAS Topic Connection Factory Properties page: name, jndiName, node and port
 - WAS Queue Factory Properties page: name, jndiName, node, persistence, priority and expiry
 - WAS Topic Factory Properties page: name, jndiName, topic name, persistence, priority and expiry
- Special cases:
 - The port value must be QUEUED or DIRECT.
 - The CCSID must be a numerical value between -2417483648 and 2417483647.
 - The persistence value must be APPLICATION_DEFINED, PERSISTENT, or NONPERSISTENT.
 - The priority value must be APPLICATION_DEFINED, or SPECIFIED.
 - The expiry value must be APPLICATION_DEFINED, UNLIMITED, or SPECIFIED.
 - On the WAS Topic Factory Properties, and the WAS Queue Factory Properties pages, the specified priority entry field must be an integer between 0 and 9, if the priority value is set to SPECIFIED .
 - On the WAS Topic Factory Properties, and the WAS Queue Factory Properties pages, the specified expiry entry field must be a value greater than 0 if expiry is set to SPECIFIED.
- Example:

```

<resources.jms:JMSProvider xmi:id="JMSProvider_2" name="WebSphere JMS Provider"
description="wasJMSProvider:description"
externalInitialContextFactory="wasJMSProvider:contextfactoryclass"
externalProviderURL="wasJMSProvider:providerURL">
<classpath>wasJMSProvider:classpath</classpath>
<factories xmi:type="resources.jms.internalmessaging:WASQueueConnectionFactory"
xmi:id="WASQueueConnectionFactory_1" name="wasQCF:name" jndiName="wasQCF:jndiName"
description="wasQCF:description" userID="wasQCF:user" password="{xor}KD4sDhwZZSos0i0="
node="wasQCF:Node">
<propertySet xmi:id="J2EEResourcePropertySet_8">
<resourceProperties xmi:id="J2EEResourceProperty_8" name="wasQCF:customName"
value="wasQCF:customValue"/>
</propertySet>
</factories>
<factories xmi:type="resources.jms.internalmessaging:WASTopicConnectionFactory"
xmi:id="WASTopicConnectionFactory_1" name="wasTCF:name" jndiName="wasTCF:jndiName"
description="wasTCF:description" userID="wasTCF:user" password="{xor}KD4sCxwZZTE+Mjo="
node="wasTCF:node" port="QUEUEED" clientID="wasTCF:clientId">
<propertySet xmi:id="J2EEResourcePropertySet_9">
<resourceProperties xmi:id="J2EEResourceProperty_9" name="wasTCF:customName"
value="wasTCF:customValue"/>
</propertySet>
</factories>
<factories xmi:type="resources.jms.internalmessaging:WASQueue" xmi:id="WASQueue_1"
name="wasQ:name" jndiName="wasQ:jndiName" description="wasQ:description"
node="wasQ:node" persistence="APPLICATION_DEFINED" priority="SPECIFIED"
specifiedPriority="1" expiry="SPECIFIED" specifiedExpiry="1">
<propertySet xmi:id="J2EEResourcePropertySet_10">
<resourceProperties xmi:id="J2EEResourceProperty_10" name="wasQ:customName"
value="wasQ:customValue"/>
</propertySet>
</factories>
<factories xmi:type="resources.jms.internalmessaging:WASTopic" xmi:id="WASTopic_1"
name="wasT:name" jndiName="wasT:jndiName" description="wasT:description"
topic="wasT:topicName" persistence="APPLICATION_DEFINED" priority="SPECIFIED"
specifiedPriority="1" expiry="SPECIFIED" specifiedExpiry="1">
<propertySet xmi:id="J2EEResourcePropertySet_11">
<resourceProperties xmi:id="J2EEResourceProperty_11" name="wasT:customName"
value="wasT:customValue"/>
</propertySet>
</factories>
<propertySet xmi:id="J2EEResourcePropertySet_12">
<resourceProperties xmi:id="J2EEResourceProperty_12" name="wasJMSProvider:customName"
value="wasJMSProvider:customValue"/>
</propertySet>
</resources.jms:JMSProvider>

```

Configuring new resource environment providers for application clients

During this task, you create new resource environment provider configurations for your application client.

To configure a new resource environment provider, perform the following steps:

1. Start the tool and open the EAR file for which you want to configure the new Java Message Service (JMS) provider. The EAR file contents display in a tree view.
2. Select from the tree the JAR file in which you want to configure the new JMS provider.
3. Expand the JAR file to view its contents.
4. Click the **Resource Environment Providers** folder. Take one of the following actions:
 - Right-click the provider folder, and click **New Provider**.
 - Click **Edit > New** on the menu bar.
5. Configure the JMS provider properties in the displayed fields.

6. Click **OK** when you finish.
7. Click **File > Save** on the menu bar to save your changes.

Resource environment provider settings for application clients

Use this page to specify resource environment entry properties.

To view this Application Client Resource Configuration Tool (ACRCT) page, click **File > Open**. After you browse for an EAR file, click **Open**. Expand the selected JAR file. Right-click **Resource Environment Providers**, and click **New**. The following fields are displayed on the **General** tab:

Name:

Specifies the administrative name for the resource environment provider.

Description:

Specifies a description of the resource environment provider for your administrative records.

Class Path:

Specifies the path to the JAR file that contains the implementation classes for the resource environment provider.

Custom Properties:

Specifies name-value pairs for setting additional properties on the object that is created at run time for this resource.

You must enter a name that is a public property on the object and a value that can be converted from a string to the type required by the set method of the property. The acceptable properties and values depend on the object that is created. Refer to the object documentation for a list of valid properties and values.

Configuring new resource environment entries for application clients

During this task, you create new resource environment entries for your client application.

1. Start the Application Client Resource Configuration Tool (ACRCT).
2. Open the EAR file for which you want to configure the new resource environment entry. The EAR file contents are in the displayed tree view.
3. Click the desired resource environment provider, and complete the following action to configure new providers:
 - Configure a new resource environment provider.
4. Expand the resource environment provider to view the **resource environment entries** folder.
5. Click the provider folder, and complete one of the following actions:
 - Right-click the folder and select **New Factory**.
 - Click **Edit > New** on the menu bar.
6. Configure the data source properties in the displayed fields.
7. Click **OK**.
8. Click **File > Save** on the menu bar to save your changes.

Resource environment entry settings for application clients

Use this page to specify resource environment entry properties.

To view this Application Client Resource Configuration Tool (ACRCT) page, click **File > Open**. After you browse for an EAR file, click **Open**. Expand the selected JAR file > **Resource Environment Providers > resource environment instance**. Right-click **Resource environment entry**, and click **New**. The following fields appear on the **General** tab:

Name:

Specifies the administrative name for the resource environment entry.

Description:

Specifies a description of the URL for your administrative records.

JNDI Name:

Specifies the Java Naming and Directory Interface (JNDI) name for the resource, including any naming subcontexts.

Use this name to link to the binding information of the platform. The binding associates the resources defined in the deployment descriptor of the module to the actual (physical) resources bound into JNDI by the platform.

Custom Properties:

Specifies name-value pairs for setting additional properties on the object that is created at run time for this resource.

You must enter a name that is a public property on the object and a value that can be converted from a string to the type required by the set method of the property. The acceptable properties and values depend on the object that is created. Refer to the object documentation for a list of valid properties and values.

Managing application clients

Perform the following tasks after deploying application clients. This task only applies to J2EE application clients.

1. Update data source and data source provider configurations.
2. Update URLs and URL provider configurations.
3. Update mail session configurations.
4. Update JMS provider, connection factories, and destination configurations.
5. Update MQ JMS provider, MQ connection factories, and MQ destination configurations.
6. Update Resource Environment Entry and Resource Environment Provider configurations.
7. (Optional) Remove application client resources.

Updating data source and data source provider configurations with the Application Client Resource Configuration Tool

During this task, you update the configuration of an existing data source or data source provider.

1. Start the tool and open the Enterprise Archive (EAR) file containing the data source or data source provider. The EAR file contents display in a tree view.
2. Select Java Archive (JAR) file from the navigation tree containing the data source or data source provider to update.

3. Expand the JAR file to view its contents until you locate the particular data source or data source provider to update. Take one of the following actions:
 - Right-click the data source object and click **Properties**.
 - Click **Edit > Properties** on the menu bar.
4. Update the properties in the displayed fields. For detailed field help, go to:
 - Data source provider properties
 - Data source properties
5. Click **OK** when finished.
6. Click **File > Save** on the menu bar to save your changes.

Updating URLs and URL provider configurations for application clients

1. Start the tool and open the Enterprise Archive (EAR) file containing the URL or URL provider. The EAR file contents are displayed in a tree view.
2. Select from the tree the Java Archive (JAR) file containing the URL or URL provider to update.
3. Expand the JAR file to view its contents.
4. Keep expanding the JAR file contents until you locate the particular URL or URL provider to update. Take one of the following actions:
 - a. Right-click the URL object and click **Properties**.
 - b. Click **Edit > Properties** on the menu bar.
5. Update the properties in the displayed fields.
6. Click **OK** when you finish.
7. Click **File > Save** on the menu bar to save your changes.

Updating mail session configurations for application clients

During this task, you update the configuration of an existing JavaMail session. You cannot update the name of the default JavaMail provider, and you cannot delete the default JavaMail provider from the navigation tree.

1. Start the tool and open the Enterprise Archive (EAR) file containing the JavaMail session. The EAR file contents are displayed in the navigation tree view.
2. Select the Java Archive (JAR) file containing the JavaMail session to update from the navigation tree.
3. Expand the JAR file to view its contents.
4. Keep expanding the JAR file contents until you locate the particular JavaMail session to update. Take one of the following actions:
 - a. Right-click the object and click **Properties**
 - b. Click **Edit > Properties** from the menu bar.
5. Update the properties in the displayed fields.
6. Click **OK** when you finish.
7. Select **File > Save** from the menu bar to save your changes.

Updating Java Message Service provider, connection factories, and destination configurations for application clients

During this task, you update the configuration of an existing Java Message Service (JMS) provider, connection factory or destination.

1. Start the tool and open the Enterprise Archive (EAR) file containing the Java Message Service (JMS) provider, connection factory, or destination. The EAR file contents display in a tree view.
2. Select the Java Archive (JAR) file containing the JMS provider, connection factory, or destination to update from the navigation tree.

3. Expand the JAR file to view its contents until you locate the particular JMS provider, connection factory, or destination to update. When you find it, do one of the following actions:
 - Right-click the provider, and click **Properties**.
 - Click **Edit > Properties** on the menu bar.
4. Update the properties in the displayed fields. For detailed field help, see:
 - JMS provider properties
 - WAS Queue connection factory properties
 - WAS Topic connection factory properties
 - WAS Queue destination properties
 - WAS Topic destination properties
5. Click **OK**.
6. Click **File > Save** to save your changes.

Updating MQ Java Message Service provider, MQ connection factories, and MQ destination configurations for application clients

During this task, you update the configuration of an existing MQ JMS provider, MQ connection factory, or MQ destination.

1. Start the Application Client Resource Configuration Tool (ACRCT).
2. Open the Enterprise Archive (EAR) file containing the MQ JMS provider, MQ connection factory, or MQ destination. The EAR file contents are displayed in the navigation tree view.
3. Select the Java Archive (JAR) file containing the MQ JMS provider, MQ connection factory, or MQ destination to update.
4. Expand the JAR file to view its contents until you locate the particular MQ product JMS provider, MQ connection factory or MQ destination that you want to update. Complete one of the following actions:
 - Right-click the appropriate object and click **Properties**.
 - Click **Edit > Properties** on the menu bar.
5. Update the properties in the displayed fields. For detailed field help, see:
 - JMS provider properties
 - MQ Queue connection factory properties
 - MQ Topic connection factory properties
 - MQ Queue destination properties
 - MQ Topic destination properties
6. Click **OK**.
7. Click **File > Save** to save your changes.

Updating resource environment entry and resource environment provider configurations for application clients

During this task, you update the configuration of an existing resource environment entry or resource environment provider.

1. Start the tool and open the Enterprise Archive (EAR) file containing the resource environment entry or resource environment provider. The EAR file contents display in a navigation tree view.
2. Select from the tree the Java Archive (JAR) file containing the resource environment entry or resource environment provider to update.
3. Expand the JAR file to view its contents until you locate the resource environment entry or resource environment provider to update. Take one of the following actions:
 - Right-click the resource environment object, and click **Properties**.
 - Click **Edit > Properties** on the menu bar.
4. Update the properties in the displayed fields. For detailed field help, see:
 - Resource environment provider properties

- Resource environment entry properties
5. Click **OK** when you finish.
 6. Click **File > Save** on the menu bar to save your changes.

Example: Configuring Resource Environment settings

The purpose of this topic is to help you configure Resource Environment settings.

- Required fields:
 - Resource Environment Provider page: name
 - Resource Environment Entry page: name, jndiName
- Example:

```
<resources.env:ResourceEnvironmentProvider xmi:id="ResourceEnvironmentProvider_1"
name="resourceEnvProvider:name" description="resourceEnvProvider:description">
<classpath>resourceEnvProvider:classpath</classpath>
<factories xmi:type="resources.env:ResourceEnvEntry" xmi:id="ResourceEnvEntry_1"
name="resourceEnvEntry:name" jndiName="resourceEnvEntry:jndiName"
description="resourceEnvEntry:description">
<propertySet xmi:id="J2EEResourcePropertySet_20">
<resourceProperties xmi:id="J2EEResourceProperty_22"
name="resourceEnvEntry:customName" value="resourceEnvEntry:customValue"/>
</propertySet>
</factories>
<propertySet xmi:id="J2EEResourcePropertySet_21">
<resourceProperties xmi:id="J2EEResourceProperty_23"
name="resourceEnvProvider:customName" value="resourceEnvProvider:customValue"/>
</propertySet>
</resources.env:ResourceEnvironmentProvider>
```

Example: Configuring resource environment custom settings for application clients

The purpose of this topic is to help you configure resource environment custom settings.

- The custom page applies to every resource type. You can specify as many custom names and values as you need.
- Example:

```
<propertySet xmi:id="J2EEResourcePropertySet_20">
<resourceProperties xmi:id="J2EEResourceProperty_22"
name="resourceEnvEntry:customName" value="resourceEnvEntry:customValue"/>
</propertySet>
```

Removing application client resources

The option to delete an item does not offer a confirmation dialog. As a safeguard, consider saving your work right before you begin this task. If you change your mind after removing an item, you can close the EAR file without saving your changes, canceling your deletion. Remember to close the EAR file immediately after the deletion, or you also lose any unsaved work that you performed since the deletion.

This task only applies to J2EE application clients.

1. Start the Application Client Resource Configuration Tool (ACRCT) and open the Enterprise Archive (EAR) file from which you want to remove an object. The EAR file contents display in the navigation tree view. If you already have an EAR file open and have made some changes, click **File > Save** to save your work before proceeding to delete an object.
2. Locate the object that you want to remove in the tree.
3. Right-click the object, and click **Delete**.
4. Click **File > Save**.

Running application clients

The J2EE specification requires support for a client container that runs stand-alone Java applications (known as J2EE application clients) and provides J2EE services to the applications. J2EE services include naming, security, and resource connections.

You are ready to run your application client using this tool after you have:

1. Written the application client program.
2. Assembled and installed an application module (.ear file) in the application server run time.
3. Deployed the application using the Application Client Resource Configuration Tool (ACRCT).

This task only applies to J2EE application clients.

1. Open a command window and invoke the following script to launch J2EE application clients using the `launchClient` shell:

```
install_root/bin/launchClient.bat
```

The `launchClient` batch command starts the application client run time, which:

- Initializes the client run time.
- Loads the class that you designated as the main class with the Assembly Toolkit.
- Runs the main method of the application client program.

When your program terminates, the application client run time cleans up the environment and the Java virtual machine (JVM) code ends.

2. Pass parameters to the `launchClient` command or to your application client program as well. The `launchClient` command allows you to do both. The `launchClient` command requires that the first parameter is either:

- An EAR file specifying the application client to launch.
- A request for `launchClient` usage information.

All other parameters intended for the `launchClient` command must begin with the `-CC` prefix.

Parameters that are not EAR files, or usage requests, or that do not begin with the `-CC` prefix, are ignored by the application client run time, and are passed directly to the application client program.

The `launchClient` command retrieves parameters from three places:

- The command line
- A properties file
- System properties

The parameters are resolved in the order listed above, with command line values having the highest priority and system properties the lowest. Using this prioritization you can set and override default values.

3. Specify the server name. By default, the **launchClient** command uses the environment variable `COMPUTERNAME` for the `BootstrapHost` property value. This setting is effective for testing your application client when it is installed on the same computer as the server. However, in other cases override this value with the name of your server.

You can override the `BootstrapHost` value by invoking `launchClient` command with the following parameters:

```
launchClient myapp.ear -CCBootstrapHost=abc.midwest.mycompany.com
```

You can also override the default by specifying the value in a properties file and passing the file name to the `launchClient` shell.

Security is controlled by the server. You do not need to configure security on the client because the client assumes that security is enabled. If security is not enabled, the server ignores the security request, and the application client works as expected.

You can store `launchClient` values in a properties file, which is a good method for distributing default values. You can then override one or more values on the command line. The format of the file is one `launchClient -CC` parameter per line without the `-CC` prefix. For example:

```
verbose=true classpath=c:\mydir\util.jar;c:\mydir\harness.jar;c:\production\G19
\global.jar BootstrapHost=abc.westcoast.mycompany.com tracefile=c:\WebSphere\mylog.txt
```

launchClient tool

This section describes the Java 2 Platform Enterprise Edition (J2EE) command line syntax for the launchClient tool for WebSphere Application Server.

The following example illustrates the command line invocation syntax for the launchClient tool:

```
launchClient [<userapp.ear>] [-help|-?] [-CCname=value] [app args]
```

where

- *userapp.ear* is the path and the name of the EAR file that contains the application client
- *name* is the name of the parameter
- *value* is the value to which the parameter ID is set
- *app args* are arguments that pass to the application client

To print the usage information, the first parameter must be a path and a name to an Enterprise Archive (EAR) file, -help, or -?. All other parameters are optional and can appear in any order. The application client run time ignores any optional parameters that do not begin with a -CC prefix, and passes them to the application client.

Parameters

Supported arguments include:

-CCsoapConnectorPort

The Simple Object Access Protocol (SOAP) connector port. If you do not specify this argument, the WebSphere Application Server default value is used.

-CCverbose

This option displays additional information messages. The default is false.

-CCclasspath

A class path value. When you launch an application, the system class path is not used. If you want to access classes that are not in the EAR file or part of the resource class paths, specify the appropriate class path here. Multiple paths can be concatenated.

-CCjar

The name of the client Java Archive (JAR) file that resides within the EAR file for the application you wish to launch. Use this argument when you have multiple client JAR files in the EAR file.

-CCadminConnectorHost

Specifies the host name of the server from which configuration information is retrieved. The default is the value of the -CCBootstrapHost parameter or the value of the local host if the -CCBootstrapHost parameter is not specified.

-CCadminConnectorPort

Indicates the port number for the administrative client function to use. The default value is 8880 for SOAP connections and 2809 for Remote Method Invocation (RMI) connections.

-CCadminConnectorType

Specifies how the administrative client connects to the server. Specify RMI to use the RMI connection type, or specify SOAP to use the SOAP connection type. The default value is SOAP.

-CCadminConnectorUser

Administrative clients use this user name when a server requires authentication. If the connection type is SOAP, and security is enabled on the server, this parameter is required. The SOAP connector does not prompt for authentication.

-CCadminConnectorPassword

The password for the user name that the `-CCadminConnectorUser` parameter specifies.

-CCaltDD

The name of an alternate deployment descriptor file. This parameter is used with the `-CCjar` parameter to specify the deployment descriptor to use. Use this argument when a client JAR file is configured with more than one deployment descriptor. Set the value to `null` to use the client JAR file standard deployment descriptor.

-CCbootstrapHost

The name of the host server you want to connect to initially. The format is:

your_server_of_choice.com

-CCbootstrapPort

The server port number. If you do not specify this argument, the WebSphere Application Server default value is used.

-CCproviderURL

Provides bootstrap server information that the initial context factory can use to obtain an initial context. WebSphere Application Server initial context factory can use either a Common Object Request Broker Architecture (CORBA) object URL or an Internet Inter-ORB Protocol (IIOP) URL. CORBA object URLs are more flexible than IIOP URLs and are the recommended URL format to use. This value can contain more than one bootstrap server address. This feature can be used when attempting to obtain an initial context from a server cluster. You can specify bootstrap server addresses, for all servers in the cluster, in the URL. The operation will succeed if at least one of the servers is running, eliminating a single point of failure. The address list does not process in a particular order. For naming operations, this value overrides the `-CCbootstrapHost` and `-CCbootstrapPort` parameters. A CORBA object URL specifying multiple systems is illustrated in the following example:

```
-CCproviderURL=corbaloc:iiop:myserver.mycompany.com:9810,:mybackupserver.mycompany.com:2809
```

This value is mapped to the `java.naming.provider.url` system property.

-CCinitonly

Use this option to initialize application client run time for ActiveX application clients without launching the client application. The default is `false`.

-CCtrace

Use this option to obtain debug trace information. You might need this information when reporting a problem to IBM customer support. The default is `false`.

-CCtracefile

Indicates the name of the file to which trace information is written. The default is to write output to the console.

-CCpropfile

Indicates the name of a properties file that contains `launchClient` properties. Specify the properties without the `-CC` prefix in the file. For example: `verbose=true`.

-CCsecurityManager

Enables and runs the WebSphere Application Server with a security manager. The default is `disable`.

-CCsecurityMgrClass

Indicates the fully qualified name of a class that implements a security manager. Only use this argument if the `-CCsecurityManager` parameter is set to `enable`. The default is `java.lang.SecurityManager`.

-CCsecurityMgrPolicy

Indicates the name of a security manager policy file. Only use this argument if the `-CCsecurityManager` parameter is set to `enable`. When you enable this parameter, the `java.security.policy` system property is set. The default is `<install_root>/properties/client.policy`.

-CCD

Use this option to have the WebSphere Application Server set the specified system property during initialization. Do not use the equals (=) character after the -CCD. For example:

-CCDcom.ibm.test.property=testvalue. You can specify multiple -CCD parameters. The general format of this parameter is -CCD<property key>=<property value>.

-CCexitVM

Use this option to have the WebSphere Application Server call the System.exit() method after the client application completes. The default is false.

-CCdumpJavaNameSpace

Prints out the Java portion of the Java Naming and Directory Interface (JNDI) name space for WebSphere Application Server. The true value uses the short format that prints out the binding name and the type of the object bound at that location. The long value uses the long format that prints out the binding name, bound object type, local object, type and string representation of the local object, for example, IORs and string values. The default value is false.

-CCtraceMode

Specifies the trace format to use for tracing. If the valid value, basic, is not specified the default is advanced. Basic tracing format is a more compact form of tracing.

The following examples demonstrate correct syntax.

On the Windows operating system:

```
launchClient c:\earfiles\myapp.ear -CCBootstrapHost=myWASServer -CCverbose=true  
app_parm1 app_parm2
```

On the UNIX operating system:

```
./launchClient.sh /usr/earfiles/myapp.ear -CCBootstrapHost=myWASServer -CCverbose=true  
app_parm1 app_parm2
```

Specifying the directory for an expanded EAR file

Each time the launchClient tool is called, it extracts the Enterprise Archive (EAR) file to a random directory name in the temporary directory on your hard drive. Then the tool sets up the thread ClassLoader to use the extracted EAR file directory and JAR files included in the Manifest.mf client Java Archive (JAR) file. In a normal J2EE Java client, these files are automatically cleaned up after the application exits. This cleanup occurs when the client container shutdown hook is called. To avoid extracting the EAR file (and removing the temporary directory) each time the launchClient tool is called, complete the following steps:

1. Specify a directory to extract the EAR file by setting the com.ibm.websphere.client.applicationclient.archivedir Java system property. If the directory does not exist or is empty, the EAR file is extracted normally. If the EAR file was previously extracted, the launchClient tool reuses the directory.
2. Delete the directory before running the launchClient tool again, if you need to update your EAR file. When you call the launchClient command, it extracts the new EAR file to the directory. If you do not delete the directory or change the system property value to point to a different directory, the launchClient tool reuses the currently extracted EAR file and does not use your changed EAR file. When specifying the com.ibm.websphere.client.applicationclient.archivedir property, make sure that the directory you specify is unique for each EAR file you use. For example, do not point the MyEar1.ear and the MyEar2.ear files to the same directory.

Example: Using a Java 2 security manager with a J2EE application client

The launchClient command provides several parameters to control the Java 2 security manager. By default the launchClient command does not enable nor run with a Java 2 security manager. To enable the Java 2 security manager, add the following parameter to your launchClient command:

```
-CCsecurityManager=enable
```

For example:

```
launchClient myear.ear -CCsecurityManager=enable
```

- When the security manager is enabled, the `launchClient` command uses by default the `java.lang.SecurityManager` class and the `<install_root>/properties/client.policy` policy file. This policy file is configured to provide the standard permissions as described in the J2EE specification for J2EE application clients and applets. If your application receives a `java.security.AccessControlException` exception, you must add additional permissions to the `client.policy` file. For more information on adding permissions, see "Configuring client.policy files" and "AccessControlException" in the Information Center.
- You can override the default security manager class by specifying the `-CCsecurityMgrClass` parameter and the default policy file using the `-CCsecurityMgrPolicy` parameter. For more information, see `launchClient` tool.
- If you invoke Java code to start the `launchClient` class, it is recommended that you do not use the `-Djava.security.manager` parameter to enable the Java 2 security manager. Using this parameter causes the Java 2 security manager to be enabled prior to initialization of the J2EE application client run-time environment. The necessary permissions are not granted and your application might receive the `java.security.AccessControlExceptions` exception.
- When the J2EE application client run time is initialized, the Enterprise Archive (EAR) file that you specified is extracted to a random subdirectory in your users temporary directory location.

Note: If the EAR file is a set of directories and subdirectories, then it is used in place and not expanded.

The J2EE application client run time sets the

`com.ibm.websphere.client.applicationclient.archivedir` system property to the directory location of the EAR file. The `client.policy` file uses this system property to inform the security manager of the location of your application client code base and to assign the configured permissions to that code base. This activity occurs when the security manager is enabled. If the security manager is enabled at the time the Java code is started, then this system property is not set, the code base is unknown, and the permissions are not granted.

- It is recommended that you enable the security manager with the J2EE application client run time. Use the following parameter: `-CCsecurityManager=enable`.

Example: Enabling Java 2 security prior to J2EE application client run time initialization

To enable the Java 2 security prior to the Java 2 Platform Enterprise Edition (J2EE) application client run time initialization, set the `com.ibm.websphere.client.applicationclient.archivedir` system property.

Perform the following steps:

1. Set the system property to the directory where the Enterprise Archive (EAR) file is expanded, for example:

```
-Dcom.ibm.websphere.client.applicationclient.archivedir=c:\myear1 (Windows systems)
```

```
-Dcom.ibm.websphere.client.applicationclient.archivedir=/usr/mrear1 (UNIX systems)
```

2. Set the `java.security.policy` system property to use the `<WAS_HOME>/properties/client.policy` file, for example:

```
-Djava.security.policy=%WAS_HOME%\properties\client.policy (Windows systems)
```

```
-Djava.security.policy=$WAS_HOME/properties/client.policy (UNIX systems)
```

- Setting the `com.ibm.websphere.client.applicationclient.archivedir` property has the following effects:

- If the directory does not exist or if it is empty, the EAR file is extracted to that directory.
- The EAR file is reused if it was previously extracted. This reuse occurs even if the EAR file specified on the command line is different.
- The security manager grants the permissions from the `client.policy` file to that directory and all its subdirectories.

- There are two types of EAR files. The first type of EAR file is a single file that contains all the enterprise application files. The second type is a set of directories and subdirectories. The following items only apply if you are using the single file form:
 - If you need to update your EAR file, delete the directory first.
 - The new EAR file is extracted to the directory the next time you run. If you do not delete the directory or change the system property value to point to a different temporary directory, the currently extracted EAR file is reused, and your changed EAR file is not used.
- When specifying the `com.ibm.websphere.client.applicationclient.archivedir` property, verify that the directory you specify is unique for each EAR file that you use. For example, do not point `MyEar1.ear` and `MyEar2.ear` files to the same directory. You must also create all directories up to, but not including, the last directory. For example, if you set the following property:
`com.ibm.websphere.client.applicationclient.archivedir=/usr/myears/myear1`

then `usr` and `myears` must exist, but `myear1` need not exist prior to running the `launchClient` class.

Application client troubleshooting tips

This section provides some debugging tips for resolving common Java 2 Platform Enterprise Edition (J2EE) application client problems. To use this troubleshooting guide, review the trace entries for one of the J2EE application client exceptions, and then locate the exception in the guide. Some of the errors in the guide are samples, and the actual error you receive can be different than what is shown here. You might find it useful to rerun the `launchClient` command specifying the `-CCverbose=true` option. This option provides additional information when the J2EE application client run time is initializing

Error: `java.lang.NoClassDefFoundError`

Explanation	This exception is thrown when Java code cannot load the specified class.
Possible causes	<ul style="list-style-type: none">• Invalid or non-existent class• Class path problem• Manifest problem

Recommended response

Check to determine if the specified class exists in a Java Archive (JAR) file within your Enterprise Archive (EAR) file. If it does, make sure the path for the class is correct. For example, if you get the exception:

```
java.lang.NoClassDefFoundError:  
WebSphereSamples.HelloEJB.HelloHome
```

verify that the HelloHome class exists in one of the JAR files in your EAR file. If it exists, verify that the path for the class is WebSphereSamples.HelloEJB.

If both the class and path are correct, then it is a class path issue. Most likely, you do not have the failing class JAR file specified in the client JAR file manifest. To verify this situation, perform the following steps:

1. Open your EAR file with the Application Assembly Tool Assembly Toolkit and click the Application Client.
2. Add the names of the other JAR files in the EAR file to the Classpath field.

This exception is generally caused by a missing Enterprise Java Beans (EJB) module name from the Classpath field.

If you have multiple JAR files to enter in the Classpath field, be sure to separate the JAR names with spaces.

If you still have the problem, you have a situation where a class is loaded from the file system instead of the EAR file. This error is difficult to debug because the offending class is not the one specified in the exception. Instead, another class is loaded from the file system before the one specified in the exception. To correct this error, review the class paths specified with the -CClasspath option and the class paths configured with the Application Client Resource Configuration Tool. Look for classes that also exist in the EAR file. You must resolve the situation where one of the classes is found on the file system instead of in the .ear file. Remove entries from the classpaths, or include the .jar files and classes in the .ear file instead of referencing them from the file system.

If you use the -CClasspath parameter or resource classpaths in the Application Client Resource Configuration Tool, and you have configured multiple JAR files or classes, verify they are separated with the correct character for your operating system. Unlike the Classpath field in the Application Assembly Tool Assembly Toolkit, these class path fields use platform-specific separator characters, usually a colon (on UNIX platforms) or a semi-colon (on Windows systems).

Note: The system class path is not used by the Application Client run time if you use the launchClient batch or shell files. In this case, the system class path would not cause this problem. However, if you load the launchClient class directly, you do have to search through the system class path as well.

Error: com.ibm.websphere.naming.CannotInstantiateObjectException: Exception occurred while attempting to get an instance of the object for the specified reference object. [Root exception is javax.naming.NameNotFoundException: xxxxxxxxxx]

Explanation

This exception occurs when you perform a lookup on an object that is not installed on the host server. Your program can look up the name in the local client Java Naming and Directory Interface (JNDI) name space, but received a NameNotFoundException exception because it is not located on the host server. One typical example is looking up an EJB component that is not installed on the host server that you access. This exception might also occur if the JNDI name you configured in your Application Client module does not match the actual JNDI name of the resource on the host server.

Possible causes

- Incorrect host server invoked
- Resource is not defined
- Resource is not installed
- Application server is not started
- Invalid JNDI configuration

Recommended response

If you are accessing the wrong host server, run the `launchClient` command again with the `-CCBootstrapHost` parameter specifying the correct host server name. If you are accessing the correct host server, use the product `dumpnamespace` command line tool to see a listing of the host server JNDI name space. If you do not see the failing object name, the resource is either not installed on the host server or the appropriate application server is not started. If you determine the resource is already installed and started, your JNDI name in your client application does not match the global JNDI name on the host server. Use the Assembly Tool to compare the JNDI bindings value of the failing object name in the client application to the JNDI bindings value of the object in the host server application. The values must match.

Error: javax.naming.ServiceUnavailableException: A communication failure occurred while attempting to obtain an initial context using the provider url: "iiop://[invalidhostname]". Make sure that the host and port information is correct and that the server identified by the provider URL is a running name server. If no port number is specified, the default port number 2809 is used. Other possible causes include the network environment or workstation network configuration. Root exception is org.omg.CORBA.INTERNAL: JORB0050E: In Profile.getAddress(), InetAddress.getByAddress[invalidhostname] threw an UnknownHostException. minor code: 4942F5B6 completed: Maybe

Explanation

This exception occurs when you specify an invalid host server name.

Possible causes

- Incorrect host server invoked
- Invalid host server name

Recommended response

Run the `launchClient` command again and specify the correct name of your host server with the `-CCBootstrapHost` parameter.

Error: javax.naming.CommunicationException: Could not obtain an initial context due to a communication failure. Since no provider URL was specified, either the bootstrap host and port of an existing ORB was used, or a new ORB instance was created and initialized with the default bootstrap host of "localhost" and the default bootstrap port of 2809. Make sure the ORB bootstrap host and port resolve to a running name server. Root exception is org.omg.CORBA.COMM_FAILURE: WRITE_ERROR_SEND_1 minor code: 49421050 completed: No

Explanation

This exception occurs when you run the `launchClient` command to a host server that does not have the Application Server started. You also receive this exception when you specify an invalid host server name. This situation might occur if you do not specify a host server name when you run the `launchClient` tool. The default behavior is for the `launchClient` tool to run to the local host, because WebSphere Application Server does not know the name of your host server. This default behavior only works when you are running the client on the same machine with WebSphere Application Server is installed.

Possible causes

- Incorrect host server invoked
- Invalid host server name
- Invalid reference to localhost
- Application server is not started
- Invalid bootstrap port

Recommended response

If you are not running to the correct host server, run the `launchClient` command again and specify the name of your host server with the `-CCBootstrapHost` parameter. Otherwise, start the Application Server on the host server and run the `launchClient` command again.

Error: javax.naming.NameNotFoundException: Name comp/env/ejb not found in context "java:"**Explanation**

This exception is thrown when the Java code cannot locate the specified name in the local JNDI name space.

Possible causes

- No binding information for the specified name
- Binding information for the specified name is incorrect
- Wrong class loader was used to load one of the program classes
- A resource reference does not include any client configuration information
- A client container on the deployment manager is trying to use enterprise extensions (not supported)

Recommended response

Open the EAR file with the Assembly Tool, and check the bindings for the failing name. Ensure this information is correct. If you are using Resource References, open the EAR file with the Application Client Resource Configuration Tool, and verify that the Resource Reference has client configuration information and the name of the Resource Reference exactly matches the JNDI name of the client configuration. If the values are correct, you might have a class loader error.

Error: java.lang.ClassCastException: Unable to load class: org.omg.stub.WebSphereSamples.HelloEJB._HelloHome_Stub at com.ibm.rmi.javax.rmi.PortableRemoteObject.narrow(portableRemoteObject.java:269)**Explanation**

This exception occurs when the application program attempts to narrow to the EJB home class and the class loaders cannot find the EJB client side bindings.

Possible causes

- The files, `*_Stub.class` and `_Tie.class`, are not in the EJB `.jar` file
- Class loader could not find the classes

Recommended response

Look at the EJB `.jar` file located in the `.ear` file and verify the class contains the Enterprise Java Beans (EJB) client side bindings. These are class files with file names that end in `_Stub` and `_Tie`. If these files are not present, then use the Application Assembly Tool Assembly Toolkit to generate the binding classes. For more information, see article [Generating deployment code for modules](#). If the binding classes are in the EJB `.jar` file, then you might have a class loader error.

**Error: WSCL0210E: The Enterprise archive file [EAR file name] could not be found.
com.ibm.websphere.client.applicationclient.ClientContainerException:
com.ibm.etools.archive.exception.OpenFailureException**

Explanation

This error occurs when the application client run time cannot read the Enterprise Archive (EAR) file.

Possible causes

The most likely cause of this error is that the system cannot find the EAR file cannot be found in the path specified on the `launchClient` command.

Recommended response

Verify that the path and file name specified on the `launchClient` command are correct. If you are running on the Windows operating system and the path and file name are correct, use a short version of the path and file name (8 character file name and 3 character extension).

The `launchClient` command appears to hang and does not return to the command line when the client application has finished.

Explanation

When running your application client using the `launchClient` command the WebSphere Application Server run time might need to display the security login dialog. To display this dialog, WebSphere Application Server run time creates an Abstract Window Toolkit (AWT) thread. When your application returns from its main method to the application client run time, the application client run time attempts to return to the operating system and end the Java virtual machine (JVM) code. However, since there is an AWT thread, the JVM code will not end until `System.exit` is called.

Possible causes

The JVM code does not end because there is an AWT thread. Java code requires that `System.exit()` be called to end AWT threads.

Recommended response

- Modify your application to call `System.exit(0)` as the last statement.
- Use the `-CCexitVM=true` parameter when you call the `launchClient` command.

For current information available from IBM Support on known problems and their resolution, see the IBM customer support page.

IBM Support has documents that can save you time gathering information needed to resolve this problem. Before opening a PMR, see the IBM customer support page.

Chapter 8. Using Web services based on Web Services for J2EE

Decide if a Web service implementation benefits your business process.

This topic introduces you to using Web services that are based on the Web Services for Java 2 platform, Enterprise Edition (J2EE) specification. WebSphere Application Server supports Web services that are developed and implemented based on Web Services for J2EE.

Use Web services when operating across a variety of platforms, including the J2EE 1.3 and non-J2EE platforms. Web services benefit your e-business solution by integrating these enterprise systems, especially systems that have developed over a long period of time.

Using Web services makes most sense if your application's clients are non-J2EE applications, unless you have J2EE applications spread across the Web. It is recommended that you use J2EE technologies if all your clients are J2EE applications because performance can decrease when you use a Web service in a J2EE exclusive environment.

Because Web services are easily applied to existing applications and information technology assets, new solutions can be deployed quickly and recomposed to address new opportunities. As Web services become more popular, the pool of services grows, promoting development of more robust models of just-in-time application and business integration over the Internet.

To use Web services applications with WebSphere Application Server:

1. Plan to use Web services. Review the Universal Description, Discovery, and Integration (UDDI), Web services gateway and Web Services Invocation Framework concepts to learn how these components can make your Web services plan more robust.
2. (Optional) Migrate existing Web services.
3. Develop Web services.
4. Assemble Web services.
5. Deploy Web services.
6. Secure Web services.
7. Tune Web services.
8. Troubleshoot Web services.

The following is an example of how a business might use Web services.

The owner of a flower shop wants to start receiving orders from customers through the Web. She starts her venture by finding wholesale flower suppliers, pricing their product, and completing contracts for future flower orders.

Using Web services, the flower shop owner can find wholesale flower suppliers. One way she finds new suppliers is to use a UDDI registry to search for potential suppliers. She chooses the suppliers and the registry sends back information on how to contact the flower distributors that meet her criteria.

The flower shop owner can request price lists from each of the suppliers by obtaining a Web Services Description Language (WSDL) file for each potential supplier. The WSDL can be downloaded from the supplier's Web page, received through email, or retrieved from the supplier's UDDI registry entry.

The WSDL describes the procedure call. When using WebSphere Application Server, the procedure call is a Java API for XML-based remote procedure call (JAX-RPC), which helps her get price lists. The WSDL file also specifies the Universal Resource Locator (URL) where the request is to be sent.

The flower shop owner now has to compare the prices she received back from each supplier, decide which suppliers she is going to do business with, and make arrangements for future orders to be filled. The ground work has been laid for the flower shop to sell merchandise through the Web by using Web services to communicate with suppliers for the best prices and complete the ordering processes. The merchandise price lists need to be published to her Web site and she needs to provide a mechanism for customers to order flowers.

The flower supplier's Web services clients are deployed on the flower shop server. When a customer makes a transaction to purchase flowers through the Web, the order is sent to the supplier through JAX-RPC. The supplier responds by sending a confirmation with the order number and shipping date. The suppliers maintain the inventory and the flower shop owner handles billing and customer order management.

Similarly, the flower shop catalog can be composed automatically from the catalogs of all the suppliers. If the supplier ships directly to the customer, the order tracking inquiries can pass directly to the supplier's order tracking system. Web services can also be used by the supplier to send invoices for orders and by the flower shop to pay the supplier's invoices. Processes that previously required forms to be filled out manually, and faxed or mailed, can now be done automatically, saving labor costs for both the flower shop and the supplier.

Using Web services is beneficial because a much larger inventory is made available to the flower shop. There is no merchandise maintenance overhead, but the flower shop can offer their customers products that they otherwise might not have. Selling flowers through the Web increases capital for the flower shop without overhead of another store or money invested into additional product.

5.1 + For a more detailed scenario, see [Web services scenario: Overview](#) which tells the story of a fictional online garden supply retailer named Plants by WebSphere and how they incorporated the Web services concept.

Web services

Web services are self-contained, modular applications that you can describe, publish, locate, and invoke over a network.

WebSphere Application Server supports Web services that are developed and implemented based on the Web Services for Java 2 platform, Enterprise Edition (J2EE) specification.

A typical Web services scenario is a business application requesting a service from a given URL using SOAP messages over a Hypertext Transport Protocol (HTTP) or Java Messaging Service (JMS) transport. The service receives the request, processes it, and returns a response. Examples of a simple Web service include weather reports or getting stock quotes. The method call is synchronous, that is, it waits until the result is available. Transaction Web services, supporting quotes, business-to-business (B2B) or business-to-client (B2C) operations include airline reservations or purchase orders.

A Web service can be the service itself or the client that accesses the service.

Web services reflect a new, service-oriented architecture approach to programming. This approach is based on the idea of building applications by discovering and implementing network-available services, or by invoking available applications to accomplish some task. Web services deliver interoperability, for example, the ability for components created in different programming languages to work together as if they were created using the same language. Web services rely on existing transport technologies, such as HTTP, and standard data encoding techniques, such as Extensible Markup Language (XML), for invoking the implementation.

The key components of a Web service are:

- Web Services Description Language (WSDL)
WSDL is the XML-based file that describes the Web service and allows the Web service request to bind to the service.
- SOAP
SOAP is the XML-based protocol that allows the Web service request to invoke the service.
- Universal Description, Discovery and Integration Protocol (UDDI)
UDDI is the registry that hosts the service broker. UDDI is similar to the Yellow Pages in a phone book.

Web services are Web applications that allow you to be more flexible in your business processes by integrating with applications that otherwise would not communicate. The inner-library loan program at your local library is a good example of the Web services concept and its evolution. The Web service concept existed even before the term; the concept exploded with the birth of the Internet. Before, you would visit your library, search the collections and check out your books. If you didn't find the book you wanted, the librarian did a search for you by computer or phone and located the book at a nearby library. The librarian ordered the book for you and you picked it up after it was delivered to your local library. By incorporating Web services applications, you can streamline your library visit. Now, you can search the local library collection and other local libraries at the same time. When other libraries provide your library with a Web service to search their collection (the service could have been provided through UDDI), your results yield their resources. Another Web service application might enable you to check the book out and get it sent to your home. Using Web services applications saves time and creates a convenience for you, as well as freeing the librarian to do other business tasks. For a more detailed scenario, see *Web services scenario: Overview* which tells the story of a fictional online garden supply retailer named Plants by WebSphere and how they incorporated the Web services concept.

Web Services for J2EE

The *Web services for Java 2 platform, Enterprise Edition (J2EE)* specification defines the programming model and run-time architecture for implementing Web services based on the Java language. Another name for Web Services for J2EE is the Java Specification Requirements (JSR) 109. The specification includes open standards for developing and implementing Web services.

WebSphere Application Server Versions 5.0.2 and 5.1 use Web Services for J2EE 1.0 as the standard for developing and implementing Web services. Web Services for J2EE 1.0 is an API include in J2EE 1.3. Before Version 5.0.2, WebSphere Application Server developed and implemented Web services based on Apache SOAP.

Web Services for J2EE focuses on eXtensible Markup Language (XML) remote procedure call (RPC) and the Java language, including representing XML-based interface definitions in the Java language; Java language definitions in XML-based definition languages, such as SOAP, and assembling.

Web Services for J2EE is the preferred platform for Web-based programming because it provides open standards allowing different types of languages, operating systems and software to communicate seamlessly through the Internet.

In order to achieve the benefits of using Web Services for J2EE, the Web services that you want to communicate with (provided by other sources), must also be based on the Java language. These other Web services can use other operating systems and languages, but the Web service itself must be based on the Java language.

For a Java application to act as Web service client, a mapping between the Web Services Description Language (WSDL) file and the Java application must exist. The mapping is defined by the Java API for XML-based RPC (JAX-RPC) specification. You can use a Java component to implement a Web service by specifying the component's interface and binding information in the WSDL file and designing the application server infrastructure to accept the service request. This entire process encompassed is based on the Web Services for J2EE specification.

Using Web Services for J2EE in WebSphere Application Server is based on J2EE 1.3. The same standards are included in J2EE 1.4.

To review the entire Web Services for J2EE specification, see [Web services: Resources for learning](#).

Java API for XML-based remote procedure call (JAX-RPC)

The *Java API for XML-based RPC (JAX-RPC)* specification enables Java language developers to develop SOAP-based interoperable and portable Web services. JAX-RPC provides core APIs for developing and deploying Web services on a Java platform and is a required part of the J2EE 1.4 platform. JAX-RPC Web services can also be developed and deployed on J2EE 1.3 containers.

WebSphere Application Server Versions 5.0.2 and 5.1 use JAX-RPC 1.0.

The JAX-RPC standard covers the programming model and bindings for using Web Services Description Language (WSDL) for Web services in the Java language.

JAX-RPC defines the mappings between the WSDL port types and the Java interfaces, as well as between Java language and eXtensible Markup Language (XML) schema types.

To learn more about JAX-RPC see [Web services: Resources for learning](#).

SOAP

SOAP is a specification for exchange of structured information in a decentralized, distributed environment. As such, it represents the main way of communication between the three key actors in a service oriented architecture (SOA): service provider, service requestor and service broker. Then main goal of its design is to be simple and extensible. A SOAP message is used to request a Web service.

WebSphere Application Server Versions 5.0.2 and 5.1 follow the standards outlined in SOAP 1.1.

SOAP was submitted to the World Wide Web Consortium (W3C) as the basis of the eXtensible Markup Language (XML) Protocol Working Group by several companies, including IBM and Lotus.

SOAP is an XML-based protocol that consists of three parts: an *envelope* that defines a framework for describing message content and process instructions, a set of *encoding rules* for expressing instances of application-defined data types, and a *convention* for representing remote procedure calls and responses.

SOAP is transport protocol-independent and can be used in combination with a variety of protocols. In Web services that are developed and implemented for use with WebSphere Application Server, SOAP is used in combination with HyperText Transport Protocol (HTTP), HTTP extension framework, and Java Messaging Service (JMS). SOAP is also operating system independent and not tied to any programming language or component technology.

Due to these characteristics, it does not matter what technology is used to implement the client, as long as the client can issue XML messages. Similarly, the service can be implemented in any language, as long as it can process XML messages. Also, both server and client sides can reside on any suitable platform.

For more information about SOAP, see [Web services: Resources for learning](#).

SOAP with Attachments API for Java

SOAP with Attachments API for Java (SAAJ) is used for SOAP messaging that works behind the scenes in the Java API for XML-based RPC (JAX-RPC) implementation. You can also use this API to directly write SOAP messaging applications rather than using JAX-RPC. SAAJ allows you to do XML messaging from the Java platform by making method calls by creating, sending and consuming XML messages over the Internet.

WebSphere Application Server Version 5.0.2 and 5.1 uses SAAJ Version 1.1.

Messages created using SAAJ follow SOAP standards. Many of the SAAJ classes and interfaces represent XML elements in a SOAP message and have the word `element` or `SOAP`, or both, in their names.

The two main types of SOAP messages are messages with attachments and messages without attachments. SAAJ provides the `SOAPMessage` class to represent a SOAP message; the `SOAPPart` class to represent the SOAP part; and the `SOAPEnvelope` interface to represent the SOAP envelope. A SOAP message can also include one or more attachment parts in addition to the SOAP part. The SOAP part must only contain XML content. If any of the message content is not in XML format, it must occur in an attachment part. SAAJ provides the `AttachmentPart` class to represent the attachment part of a SOAP message.

All SOAP messages are sent and received over a connection. When using SAAJ, the connection is represented by a `SOAPConnection` object, which goes directly from a sender to its destination. Messages sent using SAAJ are called request-response messages. The messages are sent over a `SOAPConnection` object with the method call, which sends a message (request) and blocks the request until it receives the reply (a response).

To review the entire SAAJ API, see [Web services: Resources for learning](#).

Planning to use Web services based on Web Services for J2EE

This topic discusses how to plan your use of Web services that are developed and implemented based on the Web Services for Java 2, Enterprise Edition (J2EE) specification. Read the [Web services scenario: Overview](#) which tells the story of a fictional online garden supply retailer named *Plants by WebSphere* and how they incorporated the Web services concept.

To plan to use Web services based on Web Services for J2EE:

1. Design Web services to fit your e-business solution. Consider what you want to accomplish by using Web services, how Web services fit into your current topology, applications and programming model. Decide how the Web services will process requests on the server and how the clients will manage and use the Web service.

Design your Web services for reliability, availability, manageability and security. For example, you want your Web services to process a transaction in a reasonable time at all hours of the day and provide users with good security characteristics, such as authentication for buyers. Planning to use Web services to work with WebSphere Application Server helps to meet these requirements.

To support Web services, extend WebSphere Application Server to support Web services standards. For interoperable Web services running on platforms supplied by multiple vendors, standards are essential. WebSphere Application Server uses Web services standards developed for the Java language under the Java Community Process (JCP). These standards include the Web Services for J2EE and JAX-RPC specifications.

2. Decide what development and implementation tools to use. You can use a variety of manual development and implementation tasks. Whether you have an existing Web service to implement or you want to develop your own from a Java bean or enterprise JavaBean (EJB), you can choose

different tasks respective to your resources. You can also use the WebSphere Studio Application Developer to complete development and implementation tasks.

See *Developing Web services based on Web Services for J2EE* for information about developing Web services based on the Java language through WebSphere Application Server. To read more about the WebSphere Studio Application Developer see the WebSphere Studio Application Developer Information Center.

3. Install WebSphere Application Server. (See "Installing WebSphere Application Server products" in the WebSphere Application Server Information Center.)
4. Review Web services Samples. (See "Samples Gallery" in the WebSphere Application Server Information Center.)

Develop a Web service.

Service-oriented architecture

A *service-oriented architecture (SOA)* is a collection of services that communicate with each other, for example, passing data from one service to another or coordinating an activity between one or more services.

Companies have longed to integrate existing systems in order to implement Information Technology (IT) support for business processes that cover the entire business value chain. A variety of designs are used, ranging from rigid point-to-point electronic data interchange (EDI) interactions to Web auctions. By using the Internet, companies make their IT systems available to internal departments or external customers, but the interactions are not flexible and are without standardized architecture.

Because of this increasing demand for technologies that support connecting and sharing of resources and data, there is a need for a flexible, standardized architecture. SOA is a flexible architecture that unifies business processes by structuring large applications into building blocks, or small modular functional units or services, to be used by different groups of people in and outside the company. The building blocks can be one of three roles: service provider, service broker, or service requestor. See *Web services approach to a service-oriented architecture* to learn more about these roles.

Requirements for a SOA

In order to efficiently use a SOA, you must abide by the following requirements:

- **Interoperability between different systems and programming languages .**

The most important basis for a simple integration between applications on different platforms is a communication protocol, which is available for most systems and programming languages.

- **Clear and unambiguous description language.**

To use a service offered by a provider, it is not only necessary to be able to access the provider system, but the syntax of the service interface must also be clearly defined in a platform-independent fashion.

- **Retrieval of the service.**

To allow a convenient integration at design time or even system run time, a search mechanism is required to retrieve suitable services. The services should be classified as computer-accessible, hierarchical or taxonomies based on what the services in each category do and how they can be invoked.

Web services approach to a service-oriented architecture

Web services implement a service-oriented architecture (SOA). A major focus of Web services is to make functional building blocks accessible over standard Internet protocols that are independent from platforms and programming languages. These services can be new applications or just wrapped around existing legacy systems to make them network-enabled. A service can rely on another service to achieve its goals.

Each SOA building block can play one or more of three roles:

- **Service provider**


The service provider creates a Web service and possibly publishes its interface and access information to the service registry. Each provider must decide which services to expose, how to make trade-offs between security and easy availability, how to price the services, or, if they are free, how to exploit them for other value. The provider also has to decide what category the service should be listed in for a given broker service and what sort of trading partner agreements are required to use the service.

- **Service broker**

The service broker, also known as service registry, is responsible for making the Web service interface and implementation access information available to any potential service requestor. The implementer of the broker decides about the scope of the broker. Public brokers are available through the Internet, while private brokers are only accessible to a limited audience, for example, users of a company intranet. Furthermore, the amount of the offered information has to be decided. Some brokers specialize in many listings. Others offer high levels of trust in the listed services. Some cover a broad landscape of services and others focus within an industry. There are also brokers that catalog other brokers. Depending on the business model, brokers can attempt to maximize look-up requests, number of listings or accuracy of the listings. The Universal Description, Discovery, and Integration (UDDI) specification defines a way to publish and discover information about Web services.

- **Service requestor**

The service requestor or Web service client locates entries in the broker registry using various find operations and then binds to the service provider in order to invoke one of its Web services.



Characteristics of the Web service architecture

The presented SOA employs a loose coupling between the participants, which provides greater flexibility in the following ways:

- A client is not coupled to a server, but to a service. Therefore, the integration of the server takes place outside the scope of the client application programs.
- Old and new functional blocks, or applications and systems, are encapsulated into components that work as services.
- Functional components and their interfaces are separate, allowing new interfaces to be plugged in more easily.

- Within complex applications, the control of business processes can be isolated. A business rule engine can be incorporated to control the workflow of a defined business process. Depending on the state of the workflow, the engine calls the respective services.
- Services can be incorporated dynamically during run time.
- Bindings are specified using configuration files and can be easily adapted to new needs.

Properties of a service-oriented architecture

The service-oriented architecture offers the following properties:

- **Web services are self-contained.**

On the client side, no additional software is required. A programming language with extensible markup language (XML) and Hyper Text Transport Protocol (HTTP) client support is enough to get you started. On the server side, a Web server and a SOAP server are required. It is possible to Web services-enable an existing application without writing a single line of code.

- **Web services are self-describing.**

Neither the client nor the server knows or cares about anything besides the format and content of request and response messages (loosely coupled application integration). The definition of the message format travels with the message; no external metadata repositories or code generation tool are required.

- **Web services can be published, located, and invoked across the Internet.**

This technology uses established lightweight Internet standards such as HTTP. It leverages the existing infrastructure. Some additional standards that are required to do so include SOAP, WSDL, and UDDI.

- **Web services are language-independent and interoperable.**

Client and server can be implemented in different environments. Existing code does not have to be changed in order to be Web service enabled.

- **Web services are inherently open and standard-based.**

XML and HTTP are the major technical foundation for Web services. A large part of the Web service technology has been built using open-source projects. Therefore, vendor independence and interoperability are realistic goals this time.

- **Web services are dynamic.**

Dynamic e-business can become reality using Web services because, with UDDI and WSDL, the Web service description and discovery can be automated.

- **Web services are composable.**

Simple Web services can be aggregated to more complex ones, either using workflow techniques or by calling lower-layer Web services from a Web service implementation. Web services can be chained together to perform higher-level business functions. This shortens development time and enables best-of-breed implementations.

- **Web services build on proven mature technology.**

There are a lot of commonalities, as well as a few fundamental differences to other distributed computing frameworks. For example, the transport protocol is text based and not binary.

- **Web services are loosely coupled.**

Traditionally, application design has depended on tight interconnections at both ends. Web services require a simpler level of coordination that allows a more flexible re-configuration for an integration of the services in question.

- **Web services provide programmatic access.**

The approach provides no graphical user interface; it operates at the code level. Service consumers need to know the interfaces to Web services but do not need to know the implementation details of services.

- **Web services provide the ability to wrap existing applications.**

Already existing stand-alone applications can easily be integrated into the service-oriented architecture by implementing a Web service as an interface.

Web services business models supported

The properties and benefits of using a service-oriented architecture (SOA) such as Web services is well suited for binding small modules that perform independent tasks within a highly heterogeneous e-business model. Web services can be easily wrapped around existing applications in your business model and plugged into different business processes.

For connecting to a large monolithic system that does not allow the implementation of different flexible business processes, other approaches might be better suited, for example, to satisfy specialized features, such as performance or security.

The following business models are easily implemented by using an architecture including Web services:

- **Business information**

Sharing of information with consumers or other businesses. Web services can be used to expand the reach through such services as news streams, local weather reports, integrated travel planning, and intelligent agents.

- **Business integration**

Providing transactional, fee-based services for customers. A global network of suppliers can be easily created. Web services can be implemented in auctions, e-marketplaces, and reservation systems.

- **Business process externalization**

Web services can be used to model value chains by dynamically integrating processes to a new solution within an organizational unit or even with those of other e-businesses. This can be achieved by dynamically linking internal applications to new partners and suppliers, to offer their services to complement internal services.

5.1 + To see how these models are implemented using all aspects of Web services, see Web services scenario: Overview which tells the story of a fictional online garden supply retailer named Plants by WebSphere and how they incorporate the Web services concept.

Migrating Apache SOAP Web services to Web Services for J2EE

If you have used Web services based on Apache SOAP in WebSphere Application Server Version 4.0.x through Version 5.0.2, and now want to develop and implement Web services based on the Web Services for Java 2 platform, Enterprise Edition (J2EE) specification, you need to migrate your Version 4.0 and 5.0 client applications.

To migrate these client applications according to the Web Services for J2EE standards:

1. Plan your migration strategy. There are two ways you can port an Apache SOAP client to Java API for XML-based RPC (JAX-RPC) Web services client:
 - If you have, or can create, a Web Services Description Language (WSDL) document for the service, consider using the **WSDL2Java** command tool to generate bindings for the Web service. It is more work to adapt an Apache SOAP client to use the generated JAX-RPC bindings, but the resulting client code is more robust and easier to maintain. To follow this path, see Develop a Web services client based on Web Services for J2EE.
 - If you do not have a WSDL document for the service, do not expect the service to change, and you want to port the Apache SOAP client with a minimal work, you can convert the code to use the JAX-RPC dynamic invocation interface (DII), which is similar to the Apache SOAP APIs. The DII APIs do not use WSDL or generated bindings.

You should be aware that since JAX-RPC does not specify a framework for user-written serializers, the JAX-RPC does not support the use of custom serializers. If your application cannot conform to the

default mapping between Java, WSDL, and XML supported by WebSphere Application Server, you should not attempt to migrate the application. The remainder of this topic assumes that you have decided to use the JAX-RPC DII APIs.

2. Review the GetQuote sample. (See "Samples Gallery" in the Information Center.) There is a Web services migration sample in the Samples Gallery. This sample is located in the GetQuote.java file, originally written for Apache SOAP, and includes an explanation about the changes needed to migrate to the JAX-RPC DII interfaces.
3. Convert the client application from Apache SOAP to JAX-RPC DII The Apache SOAP API and JAX-RPC DII API structures are similar. You can instantiate and configure a call object, set up the parameters, invoke the operation, and process the result in both. You can create a generic instance of a Service object with

```
javax.xml.rpc.Service service = ServiceFactory.newInstance().createService(new QName(""));
```

in JAX-RPC.

- a. Create the call object. An instance of the call object is created by

```
org.apache.soap.rpc.Call call = new org.apache.soap.rpc.Call ()
```

in Apache SOAP.

An instance of the call object is created by

```
java.xml.rpc.Call call = service.createCall();
```

in JAX-RPC.

- b. Set the endpoint URI. The target URI for the operation is passed as a parameter to

```
call.invoke: call.invoke("http://...", "");
```

in Apache SOAP.

The setTargetEndpointAddress method is used as a parameter to

```
call.setTargetEndpointAddress("http://...");
```

in JAX-RPC.

Apache SOAP has a setTargetObjectURI method on the call object that contains routing information for the request. JAX-RPC has no equivalent method. The information in the targetObjectURI is included in the targetEndpoint URI for JAX-RPC.

- c. Set the operation name. The operation name is configured on the call object by

```
call.setMethodName("opName");
```

in Apache SOAP.

The setOperationName method, which accepts a QName instead of a String parameter, is used in JAX-RPC as follows:

```
call.setOperationName(new javax.xml.namespace.Qname("namespace", "opName"));
```

- d. Set the encoding style. The encoding style is configured on the call object by

```
call.setEncodingStyleURI(org.apache.soap.Constants.NS_URI_SOAP_ENC);
```

in Apache SOAP.

The encoding style is set by a property of the call object

```
call.setProperty(javax.xml.rpc.Call.ENCODINGSTYLE_URI_PROPERTY, "http://schemas.xmlsoap.org/soap/encoding/");
```

in JAX-RPC.

- e. Declare the parameters and set the parameter values. Apache SOAP parameter types and values are described by parameter instances, which are collected into a Vector and set on the call object before the call, for example:


```

Vector params = new Vector ();
params.addElement (new org.apache.soap.rpc.Parameter(name, type, value, encodingURI));
// repeat for additional parameters...
call.setParams (params);

```

For JAX-RPC, the call object is configured with parameter names and types without providing their values, for example:

```

call.addParameter(name, xmlType, mode);
// repeat for additional parameters
call.setReturnType(type);

```

Where

- *name* (type `java.lang.String`) is the name of the parameter
- *xmlType* (type `javax.xml.namespace.QName`) is the XML type of the parameter
- *mode* (type `javax.xml.rpc.ParameterMode`) the mode of the parameter, for example, IN, OUT, or INOUT

- f. Make the call. The operation is invoked on the call object by
- ```

org.apache.soap.Response resp = call.invoke(endpointURI, "");

```

in Apache SOAP.

The parameter values are collected into an array and passed to `call.invoke` as follows:

```

Object resp = call.invoke(new Object[] {parm1, parm2,...});

```

in JAX-RPC.

- g. Check for faults. You can check for a SOAP fault on the invocation by checking the response:

```

if resp.generatedFault then {
org.apache.soap.Fault f = resp.getFault();
f.getFaultCode();
f.getFaultString();
}

```

in Apache SOAP.

A `java.rmi.RemoteException` is thrown in JAX-RPC if a SOAP fault occurs on the invocation.

```

try {
... call.invoke(...)
} catch (java.rmi.RemoteException) ...

```

- h. Retrieve the result. In Apache SOAP, if the invocation was successful and returns a result, it can be retrieved from the Response object:

```

Parameter result = resp.getReturnValue(); return result.getValue();

```

In JAX-RPC, the result of `invoke` is the returned object when no exception is thrown:

```

Object result = call.invoke(...);
...
return result;

```

Developing a Web services client based on Web Services for J2EE.

Test the Web services-enabled clients.

---

## Developing Web services based on Web Services for J2EE

This topic explains how to develop a Web service using the Web Services for Java 2, Enterprise Edition (J2EE) specification. Web services are structured in a service-oriented architecture (SOA) that makes integrating your business and e-commerce systems more flexible.

For more information about when and how you should to use Web services see Using Web services based on Web Services for J2EE. You can read about several concepts, including what is Web services,

SOAP, WSDL, Web Services for J2EE and Java API for XML-based remote procedure call (JAX-RPC). If you would like to review a scenario where Web services are used, see [Web services scenario: Overview](#).

WebSphere Application Server uses Web services standards developed for the Java language under the Java Community Process (JCP). WebSphere Application Server Versions 5.0.2 and 5.1 follow these standards:

- SOAP Version 1.1
- WSDL Version 1.1
- Web Services for J2EE (JSR-109) Version 1.0
- JAX-RPC Version 1.0
- SOAP with attachments API for Java (SAAJ) Version 1.1

You can also use the WebSphere Studio Application Developer development tools to develop Web services that integrate with WebSphere Application Server.

Before you develop a Web service you need to Set up a Web services development and unmanaged client execution environment .

Follow the [Example: Developing Web services based on Web Services for J2EE](#) for a step-by-step look at this task.

You can develop a Web service based on Web Services for J2EE in one of four ways:

1. Develop a Web service using a Java bean.
2. Develop a Web service using a stateless session enterprise bean.
3. Develop a Web service with an existing WSDL file using a Java bean.
4. Develop a Web service with an existing WSDL file using a stateless session enterprise bean.

Assemble the Web service.

## **Example: Developing Web services based on Web Services for J2EE**

This example takes you through the steps to develop a Web service from an enterprise JavaBean (EJB) implementation. The development process is based on the Web Services for Java 2 platform, Enterprise Edition (J2EE) specification.

### **1. Select the EJB or Java bean implementation that you want to enable as a Web service.**

The implementation must meet the following Web Services for J2EE specification requirements:

- It must have methods that can be mapped to a Service Endpoint Interface. See step 2 for more information.
- It must be a stateless session EJB or a Java bean without client-specific state, since the implementation bean might be selected to process a request from any client. If a client-specific state is required, a client identifier must be passed as a parameter of the Web service operation.

The selected methods of an EJB must not have a transaction attribute of Mandatory, because there is no standard for Web services transactions at this time.

A Java bean in a Web container requires the following:

- A public default constructor
- Exposed public methods
- It must not save a client-specific state between method calls
- It must be a public, non-final, and non-abstract class
- It must not define a `finalize()` method

### **2. Develop a Service Endpoint Interface.**

Developing a Web service requires a Service Endpoint Interface.

If you are using an EJB implementation, develop a Service Endpoint Interface from an EJB remote interface.

If you are using a Java bean implementation, develop a Service Endpoint Interface for a Java bean implementation.

3. Develop a WSDL file.

4. **Develop deployment descriptor templates.**

If you are using an EJB implementation, develop Web services deployment descriptor templates from an EJB implementation.

If you are using a Java bean implementation, develop Web services deployment descriptor templates for a Java bean implementation.

5. **Configure the deployment descriptors.**

By setting the `ejb-link` or `servlet-link` values of the `service-impl-bean` elements you can link to the EJB or Java bean that implements the service.

Configure the `webservices.xml` deployment descriptor.

Configure the `ibm-webservices-bnd.xmi` deployment descriptor.

6. Assemble a JAR file or Assemble a WAR file.

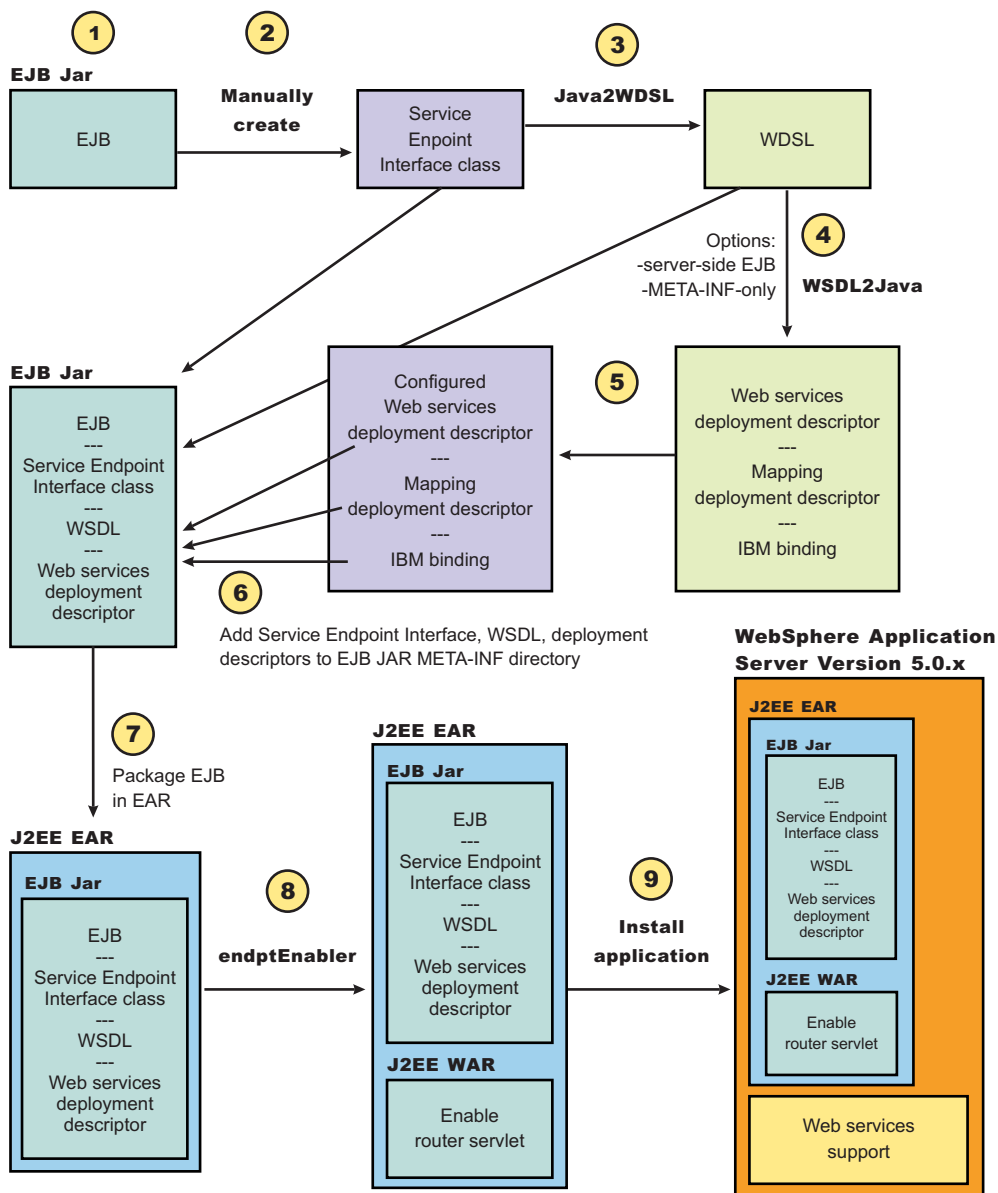
7. Assemble an EAR file from a Jar file or Assemble an EAR file from a WAR file.

8. Enable the Web service-enabled EAR file.

This step only applies if you are using an EJB implementation.

9. Deploy the Web service application.

10. Publish the WSDL file.



## Artifacts used to develop Web services based on Web Services for J2EE

*Development artifacts* enable an enterprise bean or a Java bean module to be a Web service. This topic describes artifacts used to develop Web services that are based on the Web Services for Java 2 platform, Enterprise Edition (J2EE) specification.

To create a Web service from an enterprise bean or a Java bean module, the following files are added to the respective Java archive (JAR) or Web archive (WAR) modules at assembly time:

- **Web Services Description Language (WSDL) eXtensible Markup Language (XML) file**

The WSDL XML file describes the Web service being implemented.

- **Service Endpoint Interface**

A Service Endpoint Interface is the Java interface corresponding to the Web service port type implemented. The Service Endpoint Interface is defined by the WSDL 1.1 W3C Note.

- **webservices.xml**

The `webservices.xml` file contains the J2EE Web service deployment descriptor specifying how the Web service is implemented. The `webservices.xml` file is defined in the Web Services for J2EE specification available through Web services: Resources for learning

- **ibm-webservices-bnd.xmi**

This file contains WebSphere product-specific deployment information and is defined in `ibm-webservices-bnd.xmi` assembly properties.

- **Java API for XML-based remote procedure call (JAX-RPC) mapping file**

The JAX-RPC mapping deployment descriptor specifies how Java elements are mapped to and from WSDL file elements.

The following files are added to an application client, enterprise JavaBean (EJB), or Web module to permit J2EE client access to Web services:

- **WSDL file**

The WSDL file is provided by the Web service implementer.

- **Java interfaces for the Web service**

The Java interfaces are generated from the WSDL file as specified by the JAX-RPC specification. These bindings are the Service Endpoint Interface based on the WSDL port type, or the service interface, which is based on the WSDL service.

- **webservicesclient.xml**

The `webservicesclient.xml` file is the client-side deployment descriptor describing the services being accessed. The `webservicesclient.xml` file is defined in the Web Services for J2EE specification, available through Web services: Resources for learning.

- **ibm-webservicesclient-bnd.xmi**

This file contains WebSphere product-specific deployment information such as security information.

- **Other JAX-RPC binding files**

Additional JAX-RPC binding files that support the client application in mapping SOAP to Java language are generated from WSDL by the **WSDL2Java** command tool.

## Mapping between Java language, WSDL and XML

This topic contains the mappings between the Java language and eXtensible Markup Language (XML) technologies, including XML Schema, Web Services Description Language (WSDL) and SOAP, supported by WebSphere Application Server. Most of these mappings are specified by the Java API for XML-based remote procedure call (JAX-RPC) specification. Some mappings that are optional or unspecified in JAX-RPC are also supported.

There are references to the JAX-RPC specification through this topic. You can review the JAX-RPC specification through Web services: Resources for learning.

### Notational conventions

The following table specifies the namespace prefixes and corresponding namespaces used.

| Namespace prefix | Namespace                                                                                         |
|------------------|---------------------------------------------------------------------------------------------------|
| xsd              | <a href="http://www.w3.org/2001/XMLSchema">http://www.w3.org/2001/XMLSchema</a>                   |
| xsi              | <a href="http://www.w3.org/2001/XMLSchema-instance">http://www.w3.org/2001/XMLSchema-instance</a> |
| soapenc          | <a href="http://schemas.xmlsoap.org/soap/encoding/">http://schemas.xmlsoap.org/soap/encoding/</a> |
| wSDL             | <a href="http://schemas.xmlsoap.org/wSDL/">http://schemas.xmlsoap.org/wSDL/</a>                   |
| wSDL soap        | <a href="http://schemas.xmlsoap.org/wSDL/soap/">http://schemas.xmlsoap.org/wSDL/soap/</a>         |
| ns               | user defined namespace                                                                            |
| apache           | <a href="http://xml.apache.org/xml-soap">http://xml.apache.org/xml-soap</a>                       |
| wasws            | <a href="http://websphere.ibm.com/webservices/">http://websphere.ibm.com/webservices/</a>         |

## Detailed mapping information

The following sections identify the supported mappings, including:

- Java-to-WSDL mapping
- WSDL-to-Java mapping
- Mapping between WSDL and SOAP messages

### Java-to-WSDL mapping

This section summarizes the Java-to-WSDL mapping rules. The Java-to-WSDL mapping rules are used by the **Java2WSDL** command tool for bottom-up processing. In bottom-up processing, an existing Java service implementation is used to create a WSDL file defining the Web service. The generated WSDL file can require additional manual editing for the following reasons:

- Not all Java classes and constructs have mappings to WSDL. For example, Java classes that do not comply with the Java bean specification rules might not map to a WSDL construct.
- Some Java classes and constructs have multiple mappings to WSDL. For example, a `java.lang.String` class can be mapped to either an `xsd:string` or `soapenc:string`. The **Java2WSDL** command chooses one of these mappings, but the WSDL file must be edited if a different mapping is desired.
- There are multiple ways to generate WSDL constructs. For example, the `part` element in the `wsdl:message` can be generated with a `type` or `element` attribute. The **Java2WSDL** command makes an informed choice based on the settings of the `-style` and `-use` options.
- The WSDL file describes the instance data elements sent in the SOAP message. If you want to modify the names or format used in the message, the WSDL file must be edited. For example, write a bean property value as an attribute instead of an element.
- The WSDL file requires editing if header or attachment support is desired.
- The WSDL file requires editing if a multipart WSDL, one using `wsdl:import`, is desired.

For simple services, the generated WSDL file is sufficient. For complicated services, the generated WSDL file is a good starting point.

### General issues

#### • Package to namespace mapping

The JAX-RPC specification does not specify the default mapping of Java package names to XML namespaces. The JAX-RPC specification does specify that each Java package must map to a single XML namespace. Likewise, each XML namespace must map to a single Java package. A default mapping algorithm is provided that constructs the namespace by reversing the names of the Java package and adding an `http://` prefix. For example, a package named, `com.ibm.webservice`, is mapped to the namespace `http://webservice.ibm.com`.

The default mapping between XML namespaces and Java package names can be overridden using the `-NstoPkg` and `-PkgtoNS` options of the **WSDL2Java** and **Java2WSDL** commands.

#### • Identifier mapping

Java identifiers are mapped directly to WSDL file and XML identifiers.

Java bean property names are mapped to the WSDL file and XML identifiers. For example, a Java bean, with `getInfo` and `setInfo` methods, maps to an XML construct with the name, `info`.

The Service Endpoint Interface method parameter names, if available, are mapped directly to the XML identifiers. See the **WSDL2Java** command `-implClass` option for more details.

#### • WSDL construction summary

The following table summarizes the mapping from a Java construct to the related WSDL and XML construct.

| Java construct             | WSDL and XML construct     |
|----------------------------|----------------------------|
| Service Endpoint Interface | <code>wsdl:portType</code> |

|                         |                                          |
|-------------------------|------------------------------------------|
| Method                  | wsdl:operation                           |
| Parameters              | wsdl:input, wsdl:message, wsdl:part (1)  |
| Return                  | wsdl:output, wsdl:message, wsdl:part (1) |
| Throws                  | wsdl:fault, wsdl:message, wsdl:part (1)  |
| Primitive types         | xsd and soapenc simple types             |
| Java beans              | xsd:complexType                          |
| Java bean properties    | Nested xsd:elements of xsd:complexType   |
| Arrays                  | JAX-RPC defined array xsd:complexType    |
| User defined exceptions | xsd:complexType                          |

**Note:** The generated WSDL file is affected by the `-style` and `-use` options. A `wsdl:binding` that conforms to the generated `wsdl:portType` is generated. The style and use constructs of the `wsdl:binding` are determined from the `-style` and `-use` options. A `wsdl:service` containing a port that references the generated `wsdl:binding` is generated. The names and values of the `wsdl:service` are controlled by the **Java2WSDL** command options.

- **Style and use**

Use the `-style` and `-use` options to generate different kinds of WSDL files. The four supported combinations are:

- `-style RPC -use ENCODED`
- `-style DOCUMENT -use LITERAL`
- `-style RPC -use LITERAL`
- `-style DOCUMENT -use LITERAL -wrapped false`

The `-use LITERAL` option affects the generated WSDL file in the following ways:

- The `soap:body` elements in the `wsdl:binding` are specified as `use="literal"`.
- The `soap:fault` elements in the `wsdl:binding` are specified as `use="literal"`.
- The soap encoded types are not used.
- The soap encoded array style is not used. The `maxOccurs` attribute is used to indicate arrays.

The `-use ENCODED` option affects the generated WSDL file in the following ways:

- The `soap:body` elements in the `wsdl:binding` are specified as `use="encoded"` and the `encodingStyle` is set.
- The `soap:fault` elements in the `wsdl:binding` are specified as `use="encoded"` and the `encodingStyle` is set.

The `-style RPC` option affects the generated WSDL file in the following ways:

- The `wsdl:part` elements use the `type` attribute to reference XML types.
- The `wsdl:binding` is specified as `style="rpc"`.

The `-style DOCUMENT -wrapped false` option affects the generated WSDL file in the following ways:

- The `wsdl:part` elements use the `type` attribute to reference simple types. The `element` attribute is used to reference the root `xsd:elements` for everything that is not a simple type.
- The `wsdl:binding` is specified as `style="document"`.

The `-style DOCUMENT -wrapped true` option generates a WSDL file that conforms to the .NET WSDL file format:

- A request `xsd:element` is generated for each method in the Service Endpoint Interface as follows:
  - The name of the `xsd:element` is the same as the name of the `wsdl:operation`.
  - The `xsd:element` contains an `xsd:sequence` that contains `xsd:elements` defining each parameter.
  - The request `wsdl:message` references the wrapper `xsd:element` using a single part:
    - The name of the part is `parameters`.
    - The `element` attribute is used to reference the wrapper `xsd:element`.
- A response `xsd:element` is generated for each method in the Service Endpoint Interface as follows:
  - The name of the `xsd:element` is the same as the name of the `wsdl:operation` appended with `Response`
  - The `xsd:element` contains an `xsd:sequence` that contains `xsd:elements` defining the return value.



- The request wsdl:message references this wrapper xsd:element using a single part.
  - The element attribute is used to reference the wrapper xsd:element.
- The wsdl:binding is specified as style="document".

## Mapping of standard XML types from Java types

Some Java types map directly to standard XML types. These types do not require additional XML definitions in the wsdl:types section.

### • JAX-RPC Java primitive type mapping

The following table describes the mapping from the Java primitive and standard types to XML standard types. For more information see the JAX-RPC specification.

| Java type                                                | XML type                                                                                                                                                                                                                                                                                                                                                                     |
|----------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| int                                                      | xsd:int                                                                                                                                                                                                                                                                                                                                                                      |
| long                                                     | xsd:long                                                                                                                                                                                                                                                                                                                                                                     |
| short                                                    | xsd:short                                                                                                                                                                                                                                                                                                                                                                    |
| float                                                    | xsd:float                                                                                                                                                                                                                                                                                                                                                                    |
| double                                                   | xsd:double                                                                                                                                                                                                                                                                                                                                                                   |
| boolean                                                  | xsd:boolean                                                                                                                                                                                                                                                                                                                                                                  |
| byte                                                     | xsd:byte                                                                                                                                                                                                                                                                                                                                                                     |
| byte[]                                                   | xsd:base64Binary<br><b>Note:</b> The default mapping for byte[] is xsd:base64Binary. The data in byte[] is passed over the wire as a text string encoded in the base64 format. An alternative format is xsd:hexBinary. To use the xsd:hexBinary format: <ul style="list-style-type: none"> <li>• Edit the WSDL file and change xsd:base64Binary to xsd:hexBinary.</li> </ul> |
| java.lang.String                                         | xsd:string                                                                                                                                                                                                                                                                                                                                                                   |
| java.math.BigInteger                                     | xsd:integer                                                                                                                                                                                                                                                                                                                                                                  |
| java.math.BigDecimal                                     | xsd:decimal                                                                                                                                                                                                                                                                                                                                                                  |
| java.util.Calendar                                       | xsd:dateTime                                                                                                                                                                                                                                                                                                                                                                 |
| java.util.Date                                           | xsd:date                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Note:</b> This mapping is not covered by the JAX-RPC. |                                                                                                                                                                                                                                                                                                                                                                              |
| java.lang.Boolean                                        | xsd:boolean xsi:nillable=true                                                                                                                                                                                                                                                                                                                                                |
| java.lang.Float                                          | xsd:float xsi:nillable=true                                                                                                                                                                                                                                                                                                                                                  |
| java.lang.Double                                         | xsd:double xsi:nillable=true                                                                                                                                                                                                                                                                                                                                                 |
| java.lang.Long                                           | xsd:long xsi:nillable=true                                                                                                                                                                                                                                                                                                                                                   |
| java.lang.Integer                                        | xsd:int xsi:nillable=true                                                                                                                                                                                                                                                                                                                                                    |
| java.lang.Short                                          | xsd:short xsi:nillable=true                                                                                                                                                                                                                                                                                                                                                  |
| java.lang.Byte                                           | xsd:byte xsi:nillable=true                                                                                                                                                                                                                                                                                                                                                   |

**Note:** The java.lang wrapper classes in the last seven lines of the table map to the same XML construct as the corresponding Java primitive type. In addition, the xsi:nillable attribute is generated to indicate that such elements can be null.

### • Additional Java class mappings

In addition to the standard JAX-RPC mapping, the following classes are mapped directly to XML types:

| Java type | XML type |
|-----------|----------|
|-----------|----------|

|                                                                                                                                                                              |                    |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------|
| java.util.Map<br><b>Note:</b> Any classes that implement java.util.Map are also mapped to apache:Map.                                                                        | apache:Map         |
| java.util.Collection<br><b>Note:</b> Each Java array, except byte[], and every class that implements java.util.Collection is mapped to a JAX-RPC defined soapenc:Array type. | soapenc:Array      |
| org.w3c.dom.Element                                                                                                                                                          | apache:Element     |
| java.util.Vector                                                                                                                                                             | apache:Vector      |
| java.awt.Image<br><b>Note:</b> Used for attachment support.                                                                                                                  | apache:Image       |
| javax.mail.internet.MimeMultiPart<br><b>Note:</b> Used for attachment support.                                                                                               | apache:Multipart   |
| javax.xml.transform.Source<br><b>Note:</b> Used for attachment support.                                                                                                      | apache:Source      |
| javax.activation.DataHandler<br><b>Note:</b> Used for attachment support.                                                                                                    | apache:DataHandler |

## Generation of wsdl:types

Java types that cannot be mapped directly to standard XML types are generated in the wsdl:types section.

- **Java arrays**

Java arrays for the -use ENCODED option, with the exception of byte[], are generated using the following format. See the JAX-RPC specification for more details. Alternative mappings can be found in Table 18.1 of the JAX-RPC specification.

**Java:**

```
Item[]
```

**Mapped to:**

```
<xsd:complexType name="ArrayOfItem">
 <xsd:complexContent>
 <xsd:restriction base="soapenc:Array">
 <xsd:attribute ref="soapenc:arrayType" wsdl:arrayType="ns:Item"/>
 </xsd:restriction>
 </xsd:complexContent>
</xsd:complexType>
```

- **JAX-RPC value type and bean mapping**

A Java class that matches the Java value type or bean pattern is mapped to an xsd:complexType. In order for a Java class to be mapped to XML, follow these conditions:

- The class must have a public default constructor.
- The class must not implement, directly or indirectly, java.rmi.Remote.
- Public, non-static, non-final, non-transient fields are mapped. The class can contain other fields and methods, but they are not mapped.
- If the class follows the Java bean pattern and has public getter and setter methods, the property is mapped.

Additional mapping rules can be found in the JAX-RPC specification. The following example indicates the mapping for sample base and derived Java classes:

**Java:**

```

public abstract class Base {
 public Base() {}
 public int a; // mapped
 private int b; // mapped via setter/getter
 private int c; // not mapped
 private int[] d; // mapped via indexed setter/getter

 public int getB() { return b;} // map property b
 public void setB(int b) {this.b = b;}

 public int[] getD() { return d;} // map indexed property d
 public void setD(int[] d) {this.d = d;}
 public int getD(int index) { return d[index];}
 public void setB(int index, int value) {this.d[index] = value;}

 public void someMethod() {...} // not mapped
}

public class Derived extends Base {
 public int x; // mapped
 private int y; // not mapped
}

```

#### Mapped to:

```

<xsd:complexType name="Base" abstract="true">
 <xsd:sequence>
 <xsd:element name="a" type="xsd:int"/>
 <xsd:element name="b" type="xsd:int"/>
 <xsd:element name="d" minOccurs="0" maxOccurs="unbounded" type="xsd:int"/>
 </xsd:sequence>
</xsd:complexType>

<xsd:complexType name="Derived">
 <xsd:complexContent>
 <xsd:extension base="ns:Base">
 <xsd:sequence>
 <xsd:element name="x" type="xsd:int"/>
 </xsd:sequence>
 </xsd:extension>
 </xsd:complexContent>
</xsd:complexType>

```

#### ***Inheritance and abstract classes***

The example contains two optional JAX-RPC features that are supported by WebSphere Application Server:

1. An abstract class is mapped to an `xsd:complexType` with `abstract="true"`.
2. An indexed bean property (see the methods for `d` in `Base`) are mapped to a nested element specified with `maxOccurs="unbounded"`. This format is similar to an XML array, but the SOAP message is different. An XML array defines an element for the array and nested elements for each item in the array. An element defined with `maxOccurs` indicates a series of items without the surrounding array wrapper element. Both formats are popular.

- **JAX-RPC enumeration class mapping**

Section 4.2.4 of the JAX-RPC specification defines the mapping from an XML enumeration to a Java class. Though not specifically required by the JAX-RPC, the **Java2WSDL** command performs the reverse mapping. If you have a class that has the same format as a JAX-RPC enumeration class, it is mapped to an XML enumeration.

- **Holder classes**

The JAX-RPC specification defines Holder classes in section 4.3.5. A Holder class is used to support in and out parameter passing. Every Holder class implements the `javax.xml.rpc.holders.Holder` interface. The **Java2WSDL** command maps Holder classes to the same XML type as the held type. In addition, references to Holder classes affect the generation of `wsdl:messages`.

- **Exception classes**

If a class extends the exception, `java.lang.Exception`, it is mapped to an `xsd:complexType` similar to the Java bean mapping. The getter methods of the exception are mapped as nested `xsd:elements` of the `xsd:complexType`. See section 5.5.5 of the JAX-RPC specification for more details.

**Note:** You need to generate implementation specific exception classes by invoking the **WSDL2Java** command on the resulting WSDL file.

- **Unsupported classes**

If a class cannot be mapped to an XML type, the **Java2WSDL** command issues a message and an `xsd:anyType` reference is generated in the WSDL file. In these situations, modify the Web service implementation to use the JAX-RPC compliant classes.

- **Generation of root elements**

If the **Java2WSDL** command generates an `xsd:complexType` or `xsd:simpleType` for a parameter reference, a corresponding `xsd:element` is also generated. The `xsd:element` has the same name as the `xsd:complexType/xsd:simpleType` and uses the `type` attribute to reference the `xsd:complexType/xsd:simpleType`. The `wsdl:message` part can use the `element` attribute or the `type` attribute to reference the `xsd:element` or `type`. This choice is determined by the `-style` and `-use` options.

## Generation from the interface or implementation class

The class passed to the **Java2WSDL** command represents the interface of the `wsdl:service`. The `wsdl:portType` and `wsdl:message` elements generate from this interface or implementation class.

- **Generation of the `wsdl:portType`**

The name of the `wsdl:portType` is the name of the class unless overridden by the `-portTypeName` option.

- **Generation of `wsdl:operation`**

A `wsdl:operation` generates for each public method in the interface that throws the exception, `java.rmi.RemoteException`.

- The name of the `wsdl:operation` is the name of the method.
- The `wsdl:operation` has a `parameterOrder` attribute, which defines the order of the parameters in the signature. Specifically, the `parameterOrder` lists the order of the parts of the request or response `wsdl:messages`.
- The `wsdl:operation` has a nested `wsdl:input` element that references the request `wsdl:message` using the `message` attribute.
- The `wsdl:operation` has a nested `wsdl:output` element that references the response `wsdl:message` using the `message` attribute.
- The `wsdl:operation` has a nested `wsdl:fault` element that references the default `wsdl:message` using the `message` attribute.

See sections 5.5.4 and 5.5.5 of the JAX-RPC specification for more information.

- **Generation of `wsdl:message`**

Generating the `wsdl:message` is directly related to the `-style` and `-use` options. The following is the default mapping (`-style RPC -use ENCODED`):

- A `wsdl:message` is created to represent the request.
  - A `wsdl:part` representing each parameter is added to the `wsdl:message`.
    - The `wsdl:part` has the same name as the parameter.
    - The `wsdl:part` uses the `type` attribute to locate the XML type of the parameter.
- A `wsdl:message` is created to represent the response:
  - A `wsdl:part` representing the method return is created.
    - The `wsdl:part` has the same name as the method with `Return` appended.

**Note:**

The name of the part is not specified by the JAX-RPC and is typically not checked by SOAP engines.

- The `wsdl:part` has the same name as the parameter.
- The `wsdl:part` uses the `type` attribute to locate the XML type of the parameter.

- A `wsdl:part` is created for each parameter that is a Holder.
- The `wsdl:part` has the same name as the parameter.
- A `wsdl:message` is created to represent the fault if the operation has a `wsdl:fault`.
- A `wsdl:part` representing the fault is created.
- The `wsdl:part` has the same name as the exception.
- The `wsdl:part` uses the `type` attribute to locate the `complexType` representing the exception.

The same mapping is used as described if you use the `-style RPC` and `-use LITERAL` options. It is also valid to use the `wsdl:part` element attribute instead of the `type` attribute to reference the XML schema. If you use the `-style DOCUMENT -wrapped false` and `-use LITERAL` options, the same mapping is used as described except the `wsdl:part` element attribute is used to reference the XML schema. If the XML schema is a primitive type, like `xsd:string`, the `type` attribute is used to reference the XML type. The `-style DOCUMENT, -wrapped true` and `-use LITERAL` options result in completely different mappings for the request and response messages. For example:

- A request `xsd:element` is generated for each method in the Service Endpoint Interface.
  - The name of the `xsd:element` is the same as the name of the `wsdl:operation`.
  - The `xsd:element` contains an `xsd:sequence` that contains `xsd:elements` defining each parameter.
  - The request `wsdl:message` references the wrapper `xsd:element` using a single part.
    - The name of the part is `parameters`.
    - The `element` attribute is used to reference the wrapper `xsd:element`.
- A response `xsd:element` is generated for each method in the Service Endpoint Interface.
  - The name of the `xsd:element` is the same as the name of the `wsdl:operation` appended with `Response`.
  - The `xsd:element` contains an `xsd:sequence` that contains `xsd:elements` defining the return value.
  - The request `wsdl:message` references this wrapper `xsd:element` using a single part.
    - The `element` attribute is used to reference the wrapper `xsd:element`.

- **Generation of `wsdl:binding`**

Generate a `wsdl:binding` with a name defined by the **Java2WSDL -bindingName** command.

- The `wsdlsoap:binding` style attribute is set to `rpc` if you use the `-style RPC` option; otherwise it is set to `document`.
- A `wsdl:operation` generates for each `wsdl:operation` defined in the `wsdl:portType`.
- Each `wsdl:operation` has corresponding `wsdl:input`, `wsdl:output` and `wsdl:fault` elements.
- The `wsdl:input`, `wsdl:output` and `wsdl:fault` elements each contain a `wsdlsoap:body` element.
- The `wsdlsoap:body` use attribute is set to `literal` or `encoded` according to the **-use** argument. Set the `encodingStyle` attribute to `http://schemas.xmlsoap.org/soap/encoding/` when **use** is `encoded`.

- **Generation of the `wsdl:service`**

Generate a `wsdl:service` with a name defined by the **Java2WSDL -serviceElement** command. For example:

- The `wsdl:service` contains a port with a name defined by the **Java2WSDL -servicePortName** command.
- The port references the generated `wsdl:binding` with the `binding` attribute.
- The port contains a `wsdlsoap:address` element with a
- The `location` attribute is set to the value of the **Java2WSDL -location** command.

## WSDL-to-Java mapping

The **WSDL2Java** command tool uses the following rules to generate Java classes when developing your Web services client and server. In addition, implementation specific Java classes are generated that assist in the serialization and deserialization, and invocation of the Web service.

### General issues

- **Mapping of namespace to package**

The JAX-RPC does not specify the mapping of XML namespaces to Java package names. The JAX-RPC does specify that each Java package map to a single XML namespace. Likewise, each XML namespace must map to a single Java package. A default mapping algorithm omits any protocol from

the XML namespace and reverses the names. For example, an XML namespace `http://websphere.ibm.com` becomes a Java package with the name `com.ibm.websphere`.

The default mapping of XML namespace to Java package disregards the context-root. If two namespaces are the same up until the first slash, they map to the same Java package. For example, the XML namespaces `http://websphere.ibm.com/foo` and `http://websphere.ibm.com/bar` map to the Java package `com.ibm.websphere`. The default mapping between XML namespaces and Java package names can be overridden using the `-NStoPkg` and `-PkgtoNS` options of **WSDL2Java** and **Java2WSDL** commands.

- **Identifier mapping**

XML names are much richer than Java identifiers. They can include characters that are not permitted in Java identifiers. See section 20 of the JAX-RPC specification for the rules to map an XML name to a Java identifier.

The mapping rules attempt to follow accepted Java coding conventions. Class names always begin with an uppercase letter. Method names begin with a lowercase letter. The **WSDL2Java** command generates metadata in the `_Helper` class so that the values are serialized or deserialized using the XML names specified in the WSDL file.

- **Java construction summary**

| WSDL and XML                                                                                                                                                                                                                                                             | Java                                                                                                                                                                                                                     |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>xsd:complexType (struct)</code><br><b>Note:</b> The <code>xsd:complexType</code> can also represent a Java exception if referenced by a <code>wsdl:message</code> for a <code>wsdl:fault</code> .                                                                  | Java Bean Class<br><b>Note:</b> The classes, <code>_Helper</code> , <code>_Ser</code> , and <code>_Deser</code> , generate for each Java bean class. These implementation classes aid serialization and deserialization. |
| nested <code>xsd:element/xsd:attribute</code>                                                                                                                                                                                                                            | Java bean property                                                                                                                                                                                                       |
| <code>xsd:complexType (array)</code>                                                                                                                                                                                                                                     | Java array                                                                                                                                                                                                               |
| <code>xsd:simpleType (enumeration)</code>                                                                                                                                                                                                                                | JAX-RPC enumeration class                                                                                                                                                                                                |
| <code>xsd:complexType (wrapper)</code> The method parameter signature typically is determined by the <code>wsdl:message</code> . However, if the WSDL file is a .NET wrapped style, the method parameter signature is determined by the wrapper <code>xsd:element</code> | Service Endpoint Interface method parameter signature<br><b>Note:</b> If a parameter is out or inout, a <code>Holder</code> class generates.                                                                             |
| --                                                                                                                                                                                                                                                                       | --                                                                                                                                                                                                                       |
| <code>wsdl:message</code> The method parameter signature typically is determined by the <code>wsdl:message</code> . However, if the WSDL file is a .NET wrapped style, the method parameter signature is determined by the wrapper <code>xsd:element</code>              | Service Endpoint Interface method signature<br><b>Note:</b> If a parameter is out or inout, a <code>Holder</code> class generates.                                                                                       |
| <code>wsdl:portType</code>                                                                                                                                                                                                                                               | Service Endpoint Interface                                                                                                                                                                                               |
| <code>wsdl:operation</code>                                                                                                                                                                                                                                              | Service Endpoint Interface method                                                                                                                                                                                        |
| <code>wsdl:binding</code>                                                                                                                                                                                                                                                | Stub<br><b>Note:</b> The Stub and <code>ServiceLocator</code> classes are implementation specific.                                                                                                                       |
| <code>wsdl:service</code>                                                                                                                                                                                                                                                | Service Interface and <code>ServiceLocator</code><br><b>Note:</b> The Stub and <code>ServiceLocator</code> classes are implementation specific.                                                                          |
| <code>wsdl:port</code>                                                                                                                                                                                                                                                   | Port accessor method in Service Interface                                                                                                                                                                                |

### Mapping standard XML types

- **JAX-RPC simple XML types mapping**

The following mappings are XML types to Java types. For more information about these mappings, see section 4.2.1 of the JAX-RPC specification.

| XML type                                                                                                                                                           | Java type            |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------|
| xsd:string                                                                                                                                                         | java.lang.String     |
| xsd:integer                                                                                                                                                        | java.math.BigInteger |
| xsd:int<br><b>Note:</b> If an element with this type has the xsi:nillable attribute set to true, it is mapped to the Java wrapper class of the primitive type.     | int                  |
| xsd:long<br><b>Note:</b> If an element with this type has the xsi:nillable attribute set to true, it is mapped to the Java wrapper class of the primitive type.    | long                 |
| xsd:short<br><b>Note:</b> If an element with this type has the xsi:nillable attribute set to true, it is mapped to the Java wrapper class of the primitive type.   | short                |
| xsd:decimal                                                                                                                                                        | java.math.BigDecimal |
| xsd:float<br><b>Note:</b> If an element with this type has the xsi:nillable attribute set to true, it is mapped to the Java wrapper class of the primitive type.   | float                |
| xsd:double<br><b>Note:</b> If an element with this type has the xsi:nillable attribute set to true, it is mapped to the Java wrapper class of the primitive type.  | double               |
| xsd:boolean<br><b>Note:</b> If an element with this type has the xsi:nillable attribute set to true, it is mapped to the Java wrapper class of the primitive type. | boolean              |
| xsd:byte<br><b>Note:</b> If an element with this type has the xsi:nillable attribute set to true, it is mapped to the Java wrapper class of the primitive type.    | byte                 |
| xsd:dateTime                                                                                                                                                       | java.util.Calendar   |
| xsd:date<br><b>Note:</b> This mapping is not supported by the JAX-RPC.                                                                                             | java.util.Date       |
| xsd:base64Binary                                                                                                                                                   | byte[]               |
| xsd:hexBinary                                                                                                                                                      | byte[]               |
| --                                                                                                                                                                 | --                   |
| soapenc:base64                                                                                                                                                     | byte[]               |
| soapenc:base64Binary                                                                                                                                               | byte[]               |
| soapenc:string                                                                                                                                                     | java.lang.String     |
| soapenc:boolean                                                                                                                                                    | java.lang.Boolean    |
| soapenc:float                                                                                                                                                      | java.lang.Float      |
| soapenc:double                                                                                                                                                     | java.lang.Double     |
| soapenc:decimal                                                                                                                                                    | java.math.BigDecimal |
| soapenc:int                                                                                                                                                        | java.lang.Integer    |
| soapenc:integer<br><b>Note:</b> This mapping is not supported by the JAX-RPC.                                                                                      | java.math.BigInteger |
| soapenc:short                                                                                                                                                      | java.lang.Short      |



|                                                            |                |
|------------------------------------------------------------|----------------|
| soapenc:long                                               | java.lang.Long |
| <b>Note:</b> This mapping is not supported by the JAX-RPC. |                |
| soapenc:byte                                               | java.lang.Byte |

- **JAX-RPC optional simple XML type mapping**

The **WSDL2Java** command supports the following JAX-RPC optional simple XML types.

| XML type  | Java type                 |
|-----------|---------------------------|
| xsd:qname | javax.xml.namespace.QName |

- **JAX-RPC xsd:anyType mapping**

The **WSDL2Java** command maps an `xsd:anyType` to a `java.lang.Object`. This is an optional feature of the JAX-RPC specification. The `xsd:anyType` can be used to store any XML type other than the XML primitive type. An `xsd:anyType` is always serialized along with an `xsi:type` that specifies the actual type.

- **Additional supported mappings**

The following mappings are also supported by the **WSDL2Java** command. These mappings are not defined by the JAX-RPC specification.

| XML type                           | Java type                         |
|------------------------------------|-----------------------------------|
| apache:PlainText                   | java.lang.String                  |
| <b>Note:</b> For MIME attachments. |                                   |
| apache:Map                         | java.util.Map                     |
| apache:Element                     | org.w3c.dom.Element               |
| apache:Vector                      | java.util.Vector                  |
| apache:Image                       | java.awt.Image                    |
| <b>Note:</b> For MIME attachments. |                                   |
| apache:Multipart                   | javax.mail.internet.MimeMultipart |
| <b>Note:</b> For MIME attachments. |                                   |
| apache:Source                      | javax.xml.transform.Source        |
| <b>Note:</b> For MIME attachments. |                                   |
| apache:octetStream                 | javax.activation.DataHandler      |
| <b>Note:</b> For MIME attachments. |                                   |
| apache:DataHandler                 | javax.activation.DataHandler      |
| <b>Note:</b> For MIME attachments. |                                   |

## Mapping XML defined in the `wsdl:types` section

The **WSDL2Java** command generates Java types for the XML schema constructs defined in the `wsdl:types` section. The XML schema language is broader than the required or optional subset defined by the JAX-RPC specification. The **WSDL2Java** command supports all required mappings and most optional mappings. In addition, the command supports some XML schema mappings that are outside the JAX-RPC specification. In general, the **WSDL2Java** command ignores constructs that it does not support. For example, the **WSDL2Java** command does not support the `default` attribute. If an `xsd:element` is defined with the `default` attribute, the `default` attribute is ignored. In some cases it maps unsupported constructs to `wasws:SOAPElement`.

- **Mapping of `xsd:complexType` to Java bean**

The most common mapping is from an `xsd:complexType` to a Java bean class.

- **Standard Java bean mapping**

The standard Java bean mapping is defined in section 4.2.3 of the JAX-RPC specification. The `xsd:complexType` defines the type. The nested `xsd:elements` within the `xsd:sequence` or `xsd:all` groups are mapped to Java bean properties. For example:

**XML:**

```
<xsd:complexType name="Sample">
 <xsd:sequence>
 <xsd:element name="a" type="xsd:string"/>
 <xsd:element name="b" maxOccurs="unbounded" type="xsd:string"/>
 </xsd:sequence>
</xsd:complexType>
```

**Java:**

```
public class Sample {
 // ..
 public Sample() {}

 // Bean Property a
 public String getA() {...}
 public void setA(String value) {...}

 // Indexed Bean Property b
 public String[] getB() {...}
 public String getB(int index) {...}
 public void setB(String[] values) {...}
 public void setB(int index, String value) {...}
}
}
```

– **Methods equals() and hashCode()**

The generated Java bean classes contain an implementation of the equals() method. The generation of this method is outside the JAX-RPC specification. The equals() method returns true if equals() is true for each contained bean property. The implementation accounts for self-referencing loops. This version of the equals() method is typically more useful than the "identity" equals provided by java.lang.Object.

A corresponding hashCode() method is also generated in the Java bean class.

– **Properties and indexed properties**

In the standard Java bean mapping example, the nested xsd:element for property a is mapped to a Java bean property. In addition, the **WSDL2Java** command maps a nested xsd:element with maxOccurs > 1 to a Java bean indexed property.

– **Attributes**

The **WSDL2Java** command also supports the xsd:attribute element, as shown in the following example.

Attribute a is mapped as a Java bean property, which is exactly the same mapping as a nested xsd:element. Implementation specific metadata is generated in the Sample2\_Helper class to ensure that property a is serialized and deserialized as an attribute, and not as a nested element. For example:

**XML:**

```
<xsd:complexType name="Sample2">
 <xsd:sequence>
 <xsd:attribute name="a" type="xsd:string"/>
 </xsd:sequence>
</xsd:complexType>
```

**Java:**

```
public class Sample2 {
 // ..
 public Sample2() {}
}
```

```

 // Bean Property a
 public String getA() {...}
 public void setA(String value) {...}
}

```

– **Qualified versus unqualified names**

The **WSDL2Java** command supports the `elementForm` and `attributeForm` schema attributes.

This support is not specified in the JAX-RPC specification. These attributes are used to indicate whether an element or attribute is serialized and deserialized with a qualified or unqualified name. The default setting for `elementForm` is `qualified` and the default setting for `attributeForm` is `unqualified`. These settings do not affect the Java bean class that is generated, but the information is captured in the `_Helper` class metadata.

– **Extension and the abstract attribute**

The **WSDL2Java** command supports extension of an `xsd:complexType` through the `xsd:extension` element. This support is required by the JAX-RPC specification.

The **WSDL2Java** command supports the `abstract` attribute. This feature is listed as optional by the JAX-RPC specification.

The following example shows the accepted use of the extension and abstract constructs. WebSphere Application Server uses the extension and abstract constructs to support polymorphism.

**XML:**

```

<xsd:complexType name="Base" abstract="true">
 <xsd:sequence>
 <xsd:element name="a" type="xsd:int"/>
 </xsd:sequence>
</xsd:complexType>

<xsd:complexType name="Derived">
 <xsd:complexContent>
 <xsd:extension base="ns:Base">
 <xsd:sequence>
 <xsd:element name="b" type="xsd:int"/>
 </xsd:sequence>
 </xsd:extension>
 </xsd:complexContent>
</xsd:complexType>

```

**Java:**

```

public abstract class Base {
 // ...
 public Base() {}

 public int getA() {...}
 public void setA(int a) {...}
}

public class Derived extends Base {
 // ...
 public Derived() {}

 public int getB() {...}
 public void setB(int b) {...}
}

```

– **Support for `xsd:any`**

The **WSDL2Java** command supports `xsd:anyElement`, which is different than `xsd:anyType`. This feature is not defined within the JAX-RPC specification and is subject to change.

If an `<xsd:any/>` element is defined within `xsd:sequence` or `xsd:all` group, SOAP values that do match one of the `xsd:elements` are stored in the Java bean. Values can be accessed from the Java bean using the `get_any()` and `set_any()` methods.

- **Mapping of `xsd:element`**

An `xsd:element` is a construct that has a name or name attribute, and a type defined by a `complexType` or primitive type. There are two different kinds of `xsd:elements`:

- **Root:** Defined directly underneath the schema elements and referenced by other constructs.
- **Nested:** Nested underneath group elements and are not referenced by other constructs.

Root elements are referenced by the WSDL file constructs, especially if the WSDL file is used to describe a literal service. Typically, root elements and types have the same names, which is allowed in the schema language. Under most circumstances the **WSDL2Java** command can produce Java artifacts without name collisions.

- **Four ways to reference a type**

There are four ways that a nested or root `xsd:element` can reference a type:

- **Use the type attribute:**

This is the most common way to reference a type, for example:

```
<xsd:element name="one" type="ns:myType"/>
```

The **WSDL2Java** command recognizes the type attribute as a reference to a `complexType` or `simpleTypeName`, `myType`. The **WSDL2Java** command generates a Java type based on the characteristics of `myType`. Support for the type attribute is required by the JAX-RPC specification.

- **Use the ref attribute:** For example:

```
<xsd:element ref="ns:myElement"/>
```

The **WSDL2Java** command recognizes the `ref` attribute as a reference to another root element named `myElement`. The name of the element is obtained from the referenced element, such as `myElement`. The type of the element is the type of the referenced element. The **WSDL2Java** command generates a Java type based on the characteristics of the referenced type. The `ref` attribute is an optional feature of the JAX-RPC specification.

- **Use no attribute:**

For example:

```
<xsd:element name="three"/>
```

When you do not use an attribute, the **WSDL2Java** command recognizes a reference to the `xsd:anyType` as defined by the XML schema specification. The `xsd:anyType` is an optional type of the JAX-RPC specification.

- **Use an anonymous type:**

For example:

```
<xsd:element name="four">
 <xsd:complexType>
 <xsd:sequence>
 <xsd:element name="foo" type="xsd:string"/>
 </xsd:sequence>
 </xsd:complexType>
</xsd:element>
```

When you use an anonymous type, the **WSDL2Java** command recognizes a reference to the type defined within the element.

**Note:** The `complexType` does not have a name.

The **WSDL2Java** command generates a Java type based on the characteristics of this type. Since the anonymous type does not have a name, the **WSDL2Java** command uses the name of the container element, which can result in collisions with defined types and other anonymous types.

The **WSDL2Java** command automatically detects and renames classes to avoid collisions.

Support for anonymous types is not defined by the JAX-RPC specification, however using anonymous types is common.

**Note:** An `xsd:attribute` is like an `xsd:element`; it contains a name and refers to a type. An `xsd:attribute` can refer to its type with the `typeattribute` or using an anonymous type.

– **Element specific attributes**

Some attributes can be applied to `xsd:elements` and not to XML types.

The `maxOccurs` attribute indicates the maximum number of occurrences of the element in the SOAP message. The default value is 1. If the value is greater than 1, or unbounded, the **WSDL2Java** command maps the construct to a Java array or bean indexed property. Metadata is also generated to properly serialize and deserialize a series of elements versus a normal XML array. The `maxOccurs` attribute is an optional feature of the JAX-RPC specification.

The `minOccurs` attribute indicates the minimum number of occurrences of the element in the SOAP message. The default value is 1. The `xsi:nillable` attribute indicates whether the element can have a `nil` value. The `minOccurs` and `xsi:nillable` settings affect how a `null` value is serialized in a SOAP message. If `minOccurs=0`, the `null` value is not serialized. If `xsi:nillable=true`, the value is serialized with the `xsi:nil=true` attribute.

• **Mapping of `xsd:complexType` to Java array**

The **WSDL2Java** command maps the following three kinds of XML formats to Java arrays:

**XML:**

```
<xsd:element name="array1" type="soapenc:Array"/>
```

**Java:**

```
Object[] array1;
```

**XML:**

```
<xsd:complexType name="arrayOfInt">
 <xsd:complexContent>
 <xsd:restriction base="soapenc:Array">
 <xsd:attribute ref="soapenc:arrayType" wsdl:arrayType="xsd:int[]"/>
 </xsd:restriction>
 </xsd:complexContent>
</xsd:complexType>

<xsd:element name="array2" type="ns:arrayOfInt"/>
```

**Java:**

```
int[] array2;
```

**XML:**

```
<xsd:complexType name="arrayOfInt">
 <xsd:complexContent>
 <xsd:restriction base="soapenc:Array">
 <xsd:sequence>
 <xsd:element name="item" type="xsd:int" maxOccurs="unbounded"/>
 </xsd:sequence>
 </xsd:restriction>
 </xsd:complexContent>
</xsd:complexType>

<xsd:element name="array3" type="ns:arrayOfInt"/>
```

**Java:**

```
int[] array3;
```

• **Mapping of `xsd:simpleType enumeration`**

The **WSDL2Java** command maps the following XML enumeration to a JAX-RPC specified enumeration class. See section 4.2.4 of the JAX-RPC specification for more details.

```

<xsd:simpleType name="EyeColorType">
 <xsd:restriction base="xsd:string">
 <xsd:enumeration value="brown"/>
 <xsd:enumeration value="green"/>
 <xsd:enumeration value="blue"/>
 </xsd:restriction>
</xsd:simpleType>

```

- **Mapping of xsd:complexType to exception class**

If a complexType is referenced in a wsdl:message for a wsdl:fault, the complexType is mapped to a class that extends the exception, java.lang.Exception. This mapping is similar to the mapping of a complexType to a Java bean class, except a full constructor is generated, and only getter methods are generated. See section 4.3.6 of the JAX-RPC specification for more details.

- **Other mappings**

The **WSDL2Java** command supports the mapping of xsd:simpleType and xsd:complexTypes that extend xsd:simpleTypes. These constructs are mapped to Java bean classes. The simple value is mapped to a Java bean property named, value. This is an optional feature of the JAX-RPC specification.

## Mapping of wsdl:portType

The wsdl:portType construct is mapped to the Service Endpoint Interface. The name of the wsdl:portType is mapped to the name of the class of the Service Endpoint Interface.

## Mapping of wsdl:operation

A wsdl:operation within a wsdl:portType is mapped to a method of the Service Endpoint Interface. The name of the wsdl:operation is mapped to the name of the method. The wsdl:operation contains wsdl:input and wsdl:output elements that reference the request and response wsdl:message constructs using the message attribute. The wsdl:operation can contain a wsdl:fault element that references a wsdl:message describing the fault. These faults are mapped to Java classes that extend the exception, java.lang.Exception as discussed in section 4.3.6 of the JAX-RPC specification.

- **Effect of document literal wrapped format**

If the WSDL file uses the .NET document and literal wrapped format, the method parameters are mapped from the wrapper xsd:element. The .NET document and literal format is automatically detected by the **WSDL2Java** command. The following criteria must be met:

- The WSDL file must have style="document" in its wsdl:binding constructs.
- The WSDL file must have use="literal" in its wsdl:binding constructs.
- The wsdl:message referenced by the wsdl:operation input construct must have a single part.
- The part must use the element attribute to reference an xsd:element.
- The referenced xsd:element, or wrapper element, must have the same name as the wsdl:operation.
- The wrapper element must not contain any xsd:attributes.

In such cases, each parameter name is mapped from a nested xsd:element contained within wrapper element. The type of the parameter is mapped from the type of the nested xsd:element. For example:

**XML:**

```

<xsd:element name="myMethod">
 <xsd:complexType>
 <xsd:sequence>
 <xsd:element name="param1" type="xsd:string"/>
 <xsd:element name="param2" type="xsd:int"/>
 </xsd:sequence>
 </xsd:complexType>
</xsd:element>
...
<wsdl:message name="response"/>
 <part name="parameters" element="ns:myMethod"/>
</wsdl:message name="response"/>

<wsdl:message name="response"/>

```

```

...
<wsdl:operation name="myMethod">
 <input name="input" message="request"/>
 <output name="output" message="response"/>
</wsdl:operation>

```

**Java:**

```
void myMethod(String param1, int param2) ...
```

- **Parameter mapping**

If the document and literal wrapped format is not detected, the parameter mapping follows the normal JAX-RPC mapping rules set in section 4.3.4 of the JAX-RPC specification.

Each parameter is defined by a `wsdl:message` part referenced from the input and output elements.

- A `wsdl:part` in the request `wsdl:message` is mapped to an input parameter.
- A `wsdl:part` in the response `wsdl:message` is mapped to the return value. If there are multiple `wsdl:parts` in the response message, they are mapped to output parameters.
  - A Holder class is generated for each output parameter as discussed in section 4.3.5 of the JAX-RPC specification.
- A `wsdl:part` that is both the request and response `wsdl:message` is mapped to an inout parameter.
  - A Holder class is generated for each inout parameter as discussed in section 4.3.5 of the JAX-RPC specification.
  - The `wsdl:operation parameterOrder` attribute defines the order of the parameters.

The **WSDL2Java** command supports overloaded methods, but confirm that the part names of the overloaded methods are unique. For example:

**XML:**

```

<wsdl:message name="request">
 <part name="param1" type="xsd:string"/>
 <part name="param2" type="xsd:int"/>
</wsdl:message name="request"/>

<wsdl:message name="response"/>

...
<wsdl:operation name="myMethod" parameterOrder="param1, param2">
 <input name="input" message="request"/>
 <output name="output" message="response"/>
</wsdl:operation>

```

**Java:**

```
void myMethod(String param1, int param2) ...
```

## Mapping of `wsdl:binding`

The **WSDL2Java** command uses the `wsdl:binding` information to generate an implementation specific client side stub. WebSphere Application Server uses the `wsdl:binding` information on the server side to properly deserialize the request, invoke the Web service, and serialize the response. The information in the `wsdl:binding` should not affect the generation of the Service Endpoint Interface, but it can when the document and literal wrapped format is used or when there are MIME attachments.

- **MIME attachments**

For a WSDL 1.1 compliant WSDL file, a part of an operation message, which is defined in the binding to be a MIME attachment, becomes a parameter of the type of the attachment regardless of the part declared. For example:

**XML:**

```

<wsdl:types>
 <schema ...>
 <complexType name="ArrayOfBinary">
 <restriction base="soapenc:Array">
 <attribute ref="soapenc:arrayType" wsdl:arrayType="xsd:binary[]" />
 </restriction>
 </complexType>
 </schema>
</wsdl:types>

```



```

 </restriction>
 </complexType>
</schema>
</wsdl:types>

<wsdl:message name="request">
 <part name="param1" type="ns:ArrayOfBinary"/>
</wsdl:message name="response"/>

<wsdl:message name="response"/>
 ...

<wsdl:operation name="myMethod">
 <input name="input" message="request"/>
 <output name="output" message="response"/>
</wsdl:operation>
 ...

<binding ...
<wsdl:operation name="myMethod">
 <input>
 <mime:multipartRelated>
 <mime:part>
 <mime:content part="param1" type="image/jpeg"/>
 </mime:part>
 </mime:multipartRelated>
 </input>
 ...
</wsdl:operation>

```

**Java:**

```
void myMethod(java.awt.Image param1) ...
```

The JAX-RPC requires support for the following MIME types:

| MIME type       | Java type                         |
|-----------------|-----------------------------------|
| image/gif       | java.awt.Image                    |
| image/jpeg      | java.awt.Image                    |
| text/plain      | java.lang.String                  |
| multipart/*     | javax.mail.internet.MimeMultipart |
| text/xml        | javax.xml.transform.Source        |
| application/xml | javax.xml.transform.Source        |

There are a number of problems with MIME attachments as they are defined in WSDL 1.1, including:

- The semantics of the mime:multipartRelated clause are not fully defined
- The semantics do not allow for arrays of MIME attachments

Because of these problems, several types are not specified by the JAX-RPC for MIME attachments.

These types are defined in the supported mappings previously discussed.

• **Headers**

A wsdl:binding can also define SOAP headers, for example:

**XML:**

```

<wsdl:message name="request">
 <part name="param1" type="xsd:string"/>
</wsdl:message/>

 ...

 <wsdl:message name="response"/>
 ...

 <wsdl:operation name="myMethod">

```

```

 <input name="input" message="request"/>
 <output name="output" message="response"/>
 </wsdl:operation>
 ...
<binding ...
 <wsdl:operation name="myMethod">
 <input>
 <soap:header message="request" part="param1" use="literal"/>
 </input>
 ...
 </wsdl:operation>

```

#### Java:

```
void myMethod(String param1) ...
```

This is an example of an explicit header or a header with a value determined from a method parameter. Instead of appearing in the `soap:body` SOAP message, the value of `param1` now appears in the `soap:header` SOAP message. The **WSDL2Java** command supports explicit headers and does not support implicit headers. Implicit headers have a value not determined by a parameter. For example, you could replace the `soap:header` clause in the example with:

```
<soap:header message="someOtherMsgNotAppearingInThePortType" part="someOtherPart"
use="literal"/>
```

**Note:** The **WSDL2Java** command supports explicit headers, but it is not considered good programming practice to use them. Headers are typically used for middleware logic, not business logic. Explicit headers place parameters used in business logic into the header.

### Mapping of `wsdl:service`

The `wsdl:service` element is mapped to a Generated Service interface. The Generated Service interface contains methods to access each of the ports in the `wsdl:service`. The Generated Service interface is discussed in sections 4.3.9, 4.3.10, and 4.3.11 of the JAX-RPC specification.

In addition, the `wsdl:service` element is mapped to the implementation-specific `ServiceLocator` class, which is an implementation of the Generated Service interface.

### Mapping between WSDL and SOAP messages

The WSDL file defines the format of the SOAP message that is sent over the wire. The **WSDL2Java** command and the WebSphere Application Server run time use the information in the WSDL file to confirm that the SOAP message is properly serialized and deserialized.

### Document versus RPC, literal versus encoded

If a `wsdl:binding` indicates a message is sent using an RPC format, the SOAP message contains an element defining the operation. If a `wsdl:binding` indicates the message is sent using a document format, the SOAP message does not contain the operation element.

If the `wsdl:part` is defined using the `type` attribute, the name and type of the part are used in the message. If the `wsdl:part` is defined using the `element` attribute, the name and type of the element are used in the message. The `element` attribute is not allowed by the JAX-RPC specification when `use="encoded"`.

If a `wsdl:binding` indicates a message is encoded, the values in the message are sent with `xsi:type` information. If a `wsdl:binding` indicates that a message is literal, the values in the message are typically not sent with `xsi:type` information. For example:

**WSDL:**

```

<xsd:element name="c" type="xsd:int"/>
...
<wsdl:message name="request">
 <part name="a" type="xsd:string"/>
 <part name="b" element="ns:c"/>
</wsdl:message>
...
<wsdl:operation name="method">
 <input message="request"/>
...

```

**RPC/ENCODED:**

```

<soap:body>
 <ns:method>
 <a xsi:type="xsd:string">ABC
 <element attribute is not allowed in rpc/encoded mode>
 </ns:method>
</soap:body>

```

**DOCUMENT/LITERAL:**

```

<soap:body>
 <a>ABC
 <c>123
</soap:body>

```

**DOCUMENT/LITERAL wrapped:**

```

<soap:body>
 <ns:method_wrapper>
 <a>ABC
 <c>123
 </ns:method_wrapper>
</soap:body>

```

The document and literal wrapped mode is the same as the document and literal mode. However, in the document and literal wrapped mode, there is only a single element within the body, and the element has the same name as the operation.

**Multi-ref processing**

If use=encoded, XML types that are not simpleTypes are passed in the SOAP message using the multi-ref attributes, href and id. The following example assumes that parameters one and two reference the same Java bean named, info containing fields a and b:

**Note:**

Deserialization produces a single instance of the info class for the encoded case and two instances are created for the literal case.

**RPC/ENCODED:**

```

<soap:body>
 <ns:method>
 <param1 href="#id1"/>
 <param2 href="#id2"/>
 </ns:method>
 <multiref id="id1" xsi:type="ns:info">
 <a xsi:type="xsi:string">hello<a>
 <b xsi:type="xsi:string">world
 </multiref>
</soap:body>

```

**RPC/LITERAL:**

```

<soap:body>
 <ns:method>
 <param1>
 <a>hello

```

```

 world
 </param1>
 <param2>
 <a>hello
 world
 </param2>
</ns:method>
</soap:body>

```

## XML arrays and the maxOccurs attribute

A SOAP message is affected by whether the element is defined by an XML array or using the maxOccurs attribute.

### WSDL:

```
<element name="foo" type="ns:ArrayOfString"/>
```

### Literal Instance:

```

<foo>
 <item>A</item>
 <item>B</item>
 <item>C</item>
</foo>

```

### WSDL:

```
<element name="foo" maxOccurs="unbounded" type="xsd:string"/>
```

### Literal Instance:

```

<foo>A</foo>
<foo>B</foo>
<foo>C</foo>

```

## minOccurs and nillable attributes

An element specified with minOccurs=0 that has a null value is not serialized in the SOAP message. An element specifying nillable="true" has a null value and is serialized into a SOAP message with the xsi:nil=true attribute. For example:

```
<a xsi:nil="true"/>
```

## Qualified versus unqualified

The XML schema attributeForm and elementForm attributes indicate whether the attributes and nested elements are serialized with qualified or unqualified names. If a part name is serialized, it is always serialized as an unqualified name.

## Java2WSDL command

The **Java2WSDL** command maps a Java class to a Web Services Description Language (WSDL) file by following the Java API for XML-based remote procedure call (JAX-RPC) specification. The **Java2WSDL** command accepts a Java class as input and produces a WSDL file representing the input class. If there is an existing file at the output location, it is overwritten. The WSDL file generated by the **Java2WSDL** command contains WSDL and XML schema constructs that are automatically derived from the input class. You can override these default values with command-line arguments.

### Command line syntax and arguments

The command line syntax is:

```
Java2WSDL [argument...] class
```

The following command-line arguments are supported:

## Required arguments

- **class**

Represents the fully qualified name of one of the following Java classes:

- Stateless session EJB remote interface that extends the javax.ejb.EJBObject class
- Service Endpoint Interface that extends the java.rmi.Remote class
- Java bean

The **Java2WSDL** command locates the class in CLASSPATH.

## Important arguments

- **-bindingName *name***

Specifies the name to use for the binding element. If not specified, the binding name is the portTypeName.

- **-help**

Displays the help message.

- **-helpX**

Displays the help message for extended options.

- **HelpOptions**

- **-debug**

Displays debug messages.

- **-outputImpl *impl-wsdl***

Specifies if you want an interface and implementation WSDL file emitted.

- **-locationImport *location-uri***

Specifies the location of the interface WSDL file if you use the -outputImpl argument specified.

- **-MIMEStyle**

Specifies a style representing Multipurpose Internet Mail Extensions (MIME) information. Valid arguments are:

- **Axis**
- **WSDL11** (default)

- **-soapAction**

Valid arguments are:

- **DEFAULT**

Sets the soapAction field according to deployment information.

- **NONE**

Sets the soapAction field to "".

- **OPERATION**

Sets the soapAction field to the operation name.

- **-stopClasses *parent* [, *parent*]**

If the -all argument is specified, the **Java2WSDL** command searches inherited classes and interfaces to construct the list of methods for WSDL file operations. The **Java2WSDL** command searches inherited classes and interfaces when generating extended complexTypes. The search stops when a class or interface is found within a package that begins with java or javax. The -stopClasses argument can be used to define additional classes that cause the search to stop.

- **-namespaceImpl *namespace***

Specifies the target namespace for the implementation WSDL if -outputImpl specified.

- **-voidReturn**

Valid arguments are:

- **ONEWAY**

Methods with void returns are one-way. This is the default for JMS transport.

- **TOWAY**

Methods with void returns are two-way. This the default for HTTP transport.

- **-wrapped *boolean***

Specifies if the WSDL file should be generated according to wrapped rules. This is only valid if use is literal. The option defaults to true.

- **-extraClasses *classes***

Specifies other classes that should be represented in the WSDL file.

- **-input *wsdl-uri***

Specifies the input WSDL file used to build an output WSDL file. Information from an existing WSDL file, whose name is specified in this option, is used with the input Java class to generate the desired output.

- **-implClass *impl-class***

The **Java2WSDL** command uses method parameter names to construct the WSDL file message part names. The command automatically obtains the message names from the debug information in the class. If the class is compiled without debug information, or if the class is an interface, the method parameter names are not available. In this case, you can use the **-implClass** argument to provide an alternative class from which to obtain method parameter names. The **impl-class** does not need to implement the class if the class is an interface, but it must implement the same methods as class.

- **-location *location***

Provides the location or Uniform Resource Locator (URL) of the service. Typically, this value fills automatically when the Web service deploys. Use this argument to specify the location if you want to generate a WSDL file containing a location URL without deploying. A warning displays to remind you that the generated WSDL file should not be published if the final location is not yet been determined. The name after the last slash or backslash is the name of the service port, unless the name is overridden by the **-servicePortName** argument. The service port address location attribute is assigned the specified value.

- **-namespace *targetNamespace***

Indicates the target namespace for the WSDL file being generated. See Mapping between Java, WSDL and XML for the algorithm used to obtain the default namespace.

- **-output *wsdl-uri***

Indicates the path and file name of the output WSDL file. If not specified, the default file, *class.wsdl*, is written into the current directory.

- **-PkgtoNS *package namespace***

Specifies the mapping of a Java package to a namespace. If there is a package without a namespace, the **Java2WSDL** command generates a namespace name. This argument can be repeated to specify mappings for multiple packages.

- **-portTypeName *name***

Specifies the name to use for the portType element. If not specified, the class name is used.

- **-serviceElementName *name***

Specifies the name of the service element.

- **-servicePortName *name***

Specifies the name of the service. If not specified, the service name is derived from the **-location** argument.

- **-style **RPC | DOCUMENT****

Specifies the WSDL style to use in the generated WSDL file. For more information about styles, see Mapping between Java, WSDL and XML. This argument is used with the **-use** argument.

If **RPC** is specified with **-use ENCODED**, or omitting use, a *style=rpc/use=encoded* WSDL file is generated. If **RPC** is specified with **-use LITERAL**, a *style=rpc/use=literal* WSDL file is generated. If **DOCUMENT** is specified with **-use LITERAL** or omitting use, a *style=document/use=literal* WSDL file is generated.

- **-transport **http | jms****

Generates SOAP bindings for either Hyper Text Transport Protocol (HTTP) (default) or Java Messaging Service (JMS). If **jms** is specified, the characters "jms" are appended to the WSDL file name to prevent overwriting an existing WSDL file for another transport. The transport option can only be specified once.

- **-use **LITERAL | ENCODED****

Specifies which style and use combinations are generated into the WSDL file when used with the `-style` argument. The combinations are `rpc` and `encoded`, `rpc` and `literal`, or `doc` and `literal`. For more information, see the Mapping between Java language, WSDL and XML.

- **-verbose**

Displays verbose messages.

## WSDL2Java command

The **WSDL2Java** command tool creates Java classes and deployment descriptor templates from a Web Services Description Language (WSDL) file using the Java API for XML-based remote procedure call (JAX-RPC) 1.0 specification. See Mapping between Java language, WSDL and XML for more information.

### Classes and files generated

The following kinds of classes and files are generated:

- **For each portType in the WSDL document (<wsdl:portType> element tag):**

- Service Endpoint Interface

- **For each service in the WSDL document (<wsdl:service> element tag):**

- Service Interface when the `-role develop-client` argument is specified.

- ServiceLocator when the `-role deploy-client` argument is specified.

This class is a WebSphere product-specific implementation of the service interface, and is not used directly.

- `webservices.xml` deployment descriptor template when the `-role develop-server` argument is specified

- `ibm-webservices-bnd.xmi` deployment descriptor template when the `-role develop-server` argument is specified.

- `ibm-webservices-ext.xmi` deployment descriptor template when the `-role develop-server` argument is specified.

- `wsdlfile_mapping.xml` JAX-RPC mapping file when the `-role develop-client` or `-role develop-server` is specified.

- `webservicesclient.xml` deployment descriptor template when the `-role develop-client` argument is specified.

- `ibm-webservicesclient-bnd.xmi` deployment descriptor template when the `-role develop-client` argument is specified.

- `ibm-webservicesclient-ext.xmi` deployment descriptor template when the `-role develop-client` argument is specified.

When the role is a server role, the container argument specifies which J2EE container the implementation uses. When the `-role develop-server -container ejb` arguments are specified, the `webservices.xml`, `ibm-webservices-bnd.xmi`, `ibm-webservicesclient-ext.xmi` and the mapping file are generated into the META-INF subdirectory. When the `-role develop-server -container web` arguments are specified, the files are generated into the WEB-INF directory.

- **For each binding in the WSDL file (<wsdl:binding> element tag):**

- A stub that implements the Service Endpoint Interface (deploy-client role)

- An implementation template for an enterprise bean and templates for the EJB remote interface and home interface generate when the `-role develop-server` and `-container-ejb` arguments are specified.

- An implementation template for the Java bean when the `-role develop-server` and `-container-web` arguments are specified.

- **Other classes and files:**

- A Java bean representing the structure of the type when the `-role develop-server` or `-role develop-client` arguments are specified for each `complexType` or `simpleType`.

- Three classes, `*_Ser.java`, `*_Deser.java`, and `*_Helper.java`, generate for each `complexType` to assist in converting the bean to SOAP and back when the `-role deploy-server` or `-role deploy-client` argument is specified.

- A `*Holder.java` class generates when the `-role develop-server` or `-role develop-client` arguments are specified for each out and inout parameter.



## Command line syntax

The command line syntax is:

```
WSDL2Java [arguments] WSDL-URI
```

## Required arguments

- **WSDL-URI**

Specifies the location of the input WSDL document using a Universal Resource Identifier (URI). You can also use a regular file path if the WSDL file is on the local file system.

## Important arguments

- **-container** *j2ee-container*

Indicates the J2EE container to be used. Valid arguments are:

- **client**  
Indicates client container.
- **ejb**  
Indicates enterprise JavaBean (EJB) container.
- **none**  
Indicates no container.
- **web**  
Indicates Web container.

If client is role, the default argument is **none**. If server is role, the container must be **ejb** or **web**. The same container option must be used for both development and deployment.

- **-deployScope** *argument*

Indicates how to deploy the server implementation. Valid arguments are:

- **Application**  
Uses one instance of the implementation class for all requests.
- **Request**  
Creates a new instance of the implementation class for each request.
- **Session**  
Creates a new instance of the implementation class for each session.

- **-genResolver**

Generates an `absolute-import resolver` class. The purpose of this class is to record the contents of the imported WSDL files used by the WSDL URI. This class is used by the runtime. It can also be used for future **WSDL2Java** command runs. This is desirable when the imported WSDL files are remote and can be inaccessible or slow to access. It also eliminates the possibility that a remote WSDL file might have different contents at run time than it did at development time. The generated class is named `_AbsoluteImportResolver.java`. You should compile and package this class with the other Java classes generated by the **WSDL2Java** command.

- **-help**

Displays a help message and exits.

- **-helpX**

Displays a help message for extended options and exits. The options are:

- **-all**  
Generates Java files for all types, even those that are not referenced.
- **-debug**  
Prints debugging information.
- **-fileNStoPkg** *filename*  
Specifies the file of namespace to package mappings. The default is `NStoPKG.properties`.
- **-genJava** *argument*  
Generates Java files. Valid arguments are:

- **IfNotExists**, default
- **Overwrite**
- **No**
- **-genXML** *argument*  
Generates the .xml and .xmi files. Valid arguments are:
  - **IfNotExists**, default
  - **Overwrite**
  - **No**
- **-password** *password*  
Specifies the login user password to access the WSDL URI.
- **-testCase**  
Generates the template for a JUnit test case for testing a Web service.
- **-user** *id*  
Specifies the login user name to access the WSDL URI.
- **-inputMappingFile** *mapping file*  
Specifies the file name of the Java to WSDL mapping file.
- **-NStoPkg** *namespace=package*  
By default, package names are automatically derived from the namespace strings in the WSDL file. For example, if the namespace is of the form *http://x.y.com* or *urn:x.y.com*, the corresponding package is *com.y.x*.  
You can provide your own mapping by using the **-NStoPkg** argument, which you can repeat as often as necessary, once for each unique namespace mapping. For example, if there is a namespace in the WSDL file called *urn:AddressFetcher2*, and you want files generated from the objects in this namespace to reside in the package *samples.addr*, provide the **-NStoPkg** *urn:AddressFetcher2=samples.addr* argument to the **WSDL2Java** command.
- **-output** *directory*  
Sets the root directory for emitted files.
- **-role** *j2ee role*  
Specifies the J2EE development role that identifies which files to generate. Valid arguments are:
  - **client**  
Combination of *develop-client* and *deploy-client*.
  - **deploy-client**  
Generates binding files for client deployment.
  - **deploy-server**  
Generates binding files for server deployment.
  - **develop-client** (default)  
Generates files for client development.
  - **develop-server**  
Generates files for server development.
  - **server**  
Combination of *develop-server* and *deploy-server*.
- **-timeout** *seconds*  
Specifies how long the **WSDL2Java** command should wait, in seconds, for the WSDL-URI to respond before giving up. The default is 45 seconds, -1 disables the timeout.
- **-useResolver** *resolver-class*  
Specifies an *absolute-import* resolver class to use during parsing. This class must have been created during a previous execution of the **WSDL2Java** command using the **-genResolver** option. The class must be available in CLASSPATH.
- **-verbose**  
Displays processing information, including the names of the generated files.

## Setting up a development and unmanaged client execution environment for Web services based on Web Services for J2EE

WebSphere Application Server provides command-line tools to develop Web services clients and implementations that are based on the Web Services for Java 2 platform, Enterprise Edition (J2EE) specification. WebSphere Application Server also includes the Assembly Toolkit that can be downloaded from the Web site [http://www-1.ibm.com/support/docview.wss?rs=180&context=SSEQTP&q=ASTK&uid=swg24005125&loc=en\\_US&cs=utf-8&lang=en+en](http://www-1.ibm.com/support/docview.wss?rs=180&context=SSEQTP&q=ASTK&uid=swg24005125&loc=en_US&cs=utf-8&lang=en+en). The Assembly Toolkit replaces the Application Assembly Tool (AAT). It is one of the tools available with the Application Server Toolkit product.

WebSphere Studio Application Developer also has GUI-based development tools to develop Web services that integrate with WebSphere Application Server.

Before you can set up a Web services development and unmanaged client execution environment within WebSphere Application Server, you must install WebSphere Application Server. (See "Installing WebSphere Application Server products" in the Information center.)

To set up a Web services development and unmanaged client execution environment:

1. Develop thin application client code on a server machine and run the `setupCmdLine` script.
2. Configure the path. You can add the WebSphere and Java `bin` directories to your path by typing:

On Windows platforms:

```
set PATH=%WAS_PATH%;%PATH%
```

On UNIX:

```
export PATH=$WAS_PATH:$PATH
```

Develop Web services based on Web Services for J2EE.

## Using the Java Messaging Service to transport Web services requests

WebSphere Application Server supports the use of the Java Messaging Service (JMS) to transport Web services requests, as an alternate to using HyperText Transport Protocol (HTTP). Using the JMS transport allows your Web service clients and servers to communicate through JMS queues and topics instead of HTTP connections. One-way and synchronous two-way requests are supported.

**Note:** A Web service must be implemented as an enterprise JavaBean (EJB) to be accessed through the JMS transport.

The benefits of using JMS include:

- Request and response messages are sent through reliable messaging.
- One-way requests allow clients and servers to be more loosely-coupled. For example, the server does not have to be active when the client sends the one-way request.
- One-way requests can be sent to multiple servers simultaneously through the use of a topic.

To use JMS as a transport for Web services requests:

1. Add a JMS binding and a SOAP address to the Web Services Description Language (WSDL) file. The WSDL file of a Web service must include a JMS binding and a SOAP address, which specifies a JMS endpoint URL string, in order to be accessible on the JMS transport. A JMS binding is a `wsdl:binding` element containing a `wsdlsoap:binding` element whose `transport` attribute ends in `soap/jms`, rather than the typical `soap/http` value.

In addition to the JMS binding, a `wsdl:port` element referencing the JMS binding must be included in the `wsdl:service` element within the WSDL file. The `wsdl:port` element should contain a `wsdlsoap:address` element whose `location` attribute specifies a JMS endpoint URL string.

**Note:** When you develop the Web service, a placeholder such as `file:/unspecified_location` can be used for the endpoint URL string.

2. Decide on the names and types of JMS objects that your application uses. Before your application can be installed, you need to:
  - a. Decide whether your Web service receives its requests from a queue or a topic.
  - b. Decide whether to use a secure destination (queue or topic) or a nonsecure destination.
  - c. Decide on the names for your destination, connection factory and listener port. The following list provides examples of the names that might be used for the mythical StockQuote Web service:
    - **Queue:** StockQuote\_Q (JNDI name: `jms/StockQuote_Q`)
    - **Connection factory:** StockQuote\_CF (JNDI name: `jms/StockQuote_CF`)
    - **Listener port:** StockQuoteEJB\_ListenerPort
3. Define the JMS administered objects. Once you have decided on the names and types of the JMS objects, use the administrative console or the **wsadmin** scripting interface to define the JMS objects.
4. Add the JMS endpoints to your EAR file using the **endptEnabler** command tool. You must run the **endptEnabler** command to add a JMS endpoint or router module for each Web service-enabled EJB JAR file contained in the EAR file. By default, the **endptEnabler** command adds only HTTP endpoints, but the `-transport jms` option can be used to request the addition of JMS endpoints.
5. Deploy the Web services application. During the install process you are prompted for two types of information for each Web service-enabled EJB JAR contained in your EAR file:
  - The Java Naming and Directory Interface (JNDI) name of the connection factory to be used by the EJB JAR file message driven bean (MDB) listener for sending reply messages.

If your Web service contains two-way operations, the MDB listener, defined inside the JMS endpoint added by **endptEnabler** command, needs to access a queue connection factory in order to add a reply message to the reply queue.

The MDB listener uses a resource environment reference of `java:comp/env/jms/WebServicesReplyQCF`. Therefore, during the application install process, you must provide the actual JNDI name of the queue connection factory that should be used by the MDB listener for that Web service. You might want to use the same connection factory that you defined for use by clients in step 2.

- The name of the listener port to be used by the MDB listener.

A listener port is an object used to associate a JMS connection factory with a JMS destination (queue or topic). When deployed, an MDB is configured with the correct listener port so that messages from the desired queue or topic are properly delivered to the MDB. During deployment, you can modify the name of the listener port associated with each MDB listener. The listener port name contained in the input EAR file displays as a default value. If you specify the correct listener port name to the **endptEnabler** command, perhaps through the use of properties, during step 3, you can accept the default value. Otherwise, enter the correct listener port name.

**Hint:** By default, the **endptEnabler** command produces listener port names of the form `<ejb-jar-name>_ListenerPort`. To simplify this step, define the listener ports that follow this naming convention during step 2.

## Java Messaging Service endpoint URL syntax

A Java Messaging Service (JMS) endpoint URL is used to access a Web service with the JMS transport. This URL specifies the JMS destination and connection factory, as well as the port component name for the Web service request. This is similar to the HTTP endpoint URL, which specifies the host and port, as well as the context root and port component name.

A JMS endpoint URL has the following general form:

```
jms:[queue|topic]?<property>=<value>&<property>=<value>&...
```

The URL consists of the transport type, `jms:`, followed by either `/queue` or `/topic` to indicate the JMS destination type, followed by the query string containing a list of property and value pairs used to specify the JMS endpoint information.

The properties supported in the URL string are described as follows:

#### Destination-related properties (required)

Property name	Description
destination	Specifies the Java Naming and Directory Interface (JNDI) name of the destination queue or topic.
connectionFactory	Specifies the JNDI name of the connection factory.
targetService	Specifies the name of the port component to which the request is dispatched.

#### JNDI-related properties (optional)

Property name	Description
initialContextFactory	Specifies the name of the initial context factory to use which is mapped to the <code>java.naming.factory.initial</code> property.
jndiProviderURL	Specifies the JNDI provider URL which is mapped to the <code>java.naming.provider.url</code> property.

#### JMS-related properties (optional)

Property name	Description
deliveryMode	Indicates whether the request message should be persistent or not. The valid values are 1 for nonpersistent and 2 for persistent. The default value is 1.
timeToLive	Specifies the lifetime, in milliseconds, of the request message. A value of 0 indicates an infinite lifetime.
priority	Specifies the JMS priority associated with the request message. Valid values are between 0 to 9. The default value is 4.

The required properties, `destination`, `connectionFactory`, and `targetService`, must appear in the JMS endpoint URL string. The rest of the properties are optional.

You can set any of the properties on the client Stub object. This means that the various properties can be specified by including them as part of the endpoint URL or they can be set programmatically by the client on the Stub object. Properties specified on the client Stub object take precedence over properties specified as part of a JMS endpoint URL string.

## Developing a Web service from a Java bean

Set up a Web services development and unmanaged client execution environment.

To develop a Web service from a Java bean:

1. Access an existing Java bean Web archive (WAR) file.
2. Develop a Java bean Service Endpoint Interface.
3. Develop a Web Services Description Language (WSDL) file.
4. Develop Web services deployment descriptor templates for a Java bean implementation.
5. Configure the `webservices.xml` deployment descriptor.

6. Configure the `ibm-webservices-bnd.xml` deployment descriptor.
7. Assemble a Web services-enabled WAR file when starting from Java.
8. Assemble a Web services-enabled WAR into an EAR file.
9. Deploy the EAR file into WebSphere Application Server.

Test the Web service.

## Developing a WSDL file

Develop a Service Endpoint Interface.

You need a Web Services Description Language (WSDL) file to use Web services. You can develop your own WSDL file or get one from a Web service provider through E-mail, downloading or through a Uniform Resource Locator (URL). This documentation assumes you are creating your own.

To develop a WSDL file:

1. Configure the Service Endpoint Interface class and referenced classes into your CLASSPATH.
  - On Windows, set `CLASSPATH=%CLASSPATH%;<list your application JAR files and classes>`.
  - On UNIX, export `CLASSPATH=%CLASSPATH:<list your application JAR files and classes>`.
2. Run the **Java2WSDL *seiInterface*** command. A WSDL file named `seiInterface.wsdl` is created.
  - Move the WSDL file to the `META-INF/wsdl` subdirectory if you are using an enterprise JavaBean (EJB).
  - Move the WSDL file to the `WEB-INF/wsdl` subdirectory if you are using a Java bean.
3. Edit the generated WSDL file and inspect the part names. The WSDL parts have names like `arg_0_0`. Modify the WSDL file to use the actual names of the Java parameters.
4. (Optional) Use the **Java2WSDL** command tool to generate the correct part names of WSDL file. You can automatically generate and set the correct part names by using the **Java2WSDL** command tool. Generating and setting the part names is done by providing additional information to the **Java2WSDL** command tool in the form of a Java implementation class that implements the same methods as the Service Endpoint Interface and is compiled with debug information on (**javac -g**). Parameter names are stored in the `.class` file with the debug information. If your implementation class was compiled with debug on, you can use the **Java2WSDL -implClass *seiImpl seiInterface*** command to generate a WSDL file having the proper part names.

A WSDL file that defines the Web service described by the Service Endpoint Interface.

This example uses a JAR file named `AddressBook.jar` containing a class named `AddressBook.class` file.

You must add the `AddressBook.jar` file to your CLASSPATH to create the WSDL file. The JAR file contains an EJB implementation class that was compiled with debugging information on. Run the **Java2WSDL -implClass *addr.AddressBookBean addr.AddressBook*** command to create a WSDL file named `AddressBook.wsdl`.

Develop Web services deployment descriptor templates from a WSDL file.

### **WSDL:**

*Web Services Description Language (WSDL)* is an eXtensible Markup Language (XML)-based description language that has been submitted to the World Wide Web Consortium (W3C) as the industry standard for describing Web services. The power of WSDL is derived from two main architectural principles: the ability to describe a set of business operations and the ability to separate the description into two basic units, a description of the operations and the details of how the operation and the information associated with it are packaged.



The WSDL document is the engine of a Java 2 platform, Enterprise Edition (J2EE) Web service; without it there is no service. The information within a WSDL file maps to the Java application to create a Web service. WebSphere Application Server Versions 5.0.2 and 5.1 use standards based on WSDL 1.1.

A WSDL document allows a service provider to specify the name and address of the Web service; protocol and encoding style used when accessing the public operations of the Web service; and the type information, including name, operations, parameters and data comprising the interface of the Web service.

A WSDL document defines services as collections of network endpoints, or ports. In WSDL, the abstract definition of endpoints and messages is separated from their concrete network deployment or data format bindings. This allows the reuse of abstract definitions: messages, which are abstract descriptions of the data being exchanged, and port types which are abstract collections of operations. The concrete protocol and data format specifications for a particular port type constitutes a reusable binding. A port is defined by associating a network address with a reusable binding, and a collection of ports define a service. Therefore, a WSDL document is composed of several elements. See WSDL anatomy for more information and examples of the WSDL elements.

When creating a Web service for WebSphere Application Server, you must first have an implementation bean that includes a Service Endpoint Interface. Then, you use the **Java2WSDL** command-line tool to create a WSDL that defines the Web service. To learn more about how the WSDL file is used in the development process, see *Developing Web services based on Web Services for J2EE*.

#### ***Publishing WSDL files:***

To publish a Web Services Description Language (WSDL) file you need an enterprise application, also known as an enterprise archive (EAR) file, that contains a Web services-enabled module and has been deployed into WebSphere Application Server. See *Deploying Web services based on Web Services for Java 2 platform, Enterprise Edition (J2EE)*.

The WSDL files for each Web services-enabled module are published to the file system location you specify. You can provide these WSDL files to clients that want to invoke your Web services.

By publishing your WSDL files, you are enabling users to generate and export WSDL files with endpoint information for Web services ports. If you are a client developer or a system administrator, you can use WSDL files to enable clients to connect to a Web service. Server configuration allows you to specify URL fragments to enable full URL specification of WSDL file ports.

You can publish WSDL files for the deployed EAR file in one of three ways:

1. Publish a WSDL file with the administrative console.
2. Publish a WSDL file with the **wsadmin** command tool.
3. Publish a WSDL file through a URL.

#### ***Publishing WSDL files with the administrative console:***

**5.1** You can specify default or custom HTTP URL prefixes when using the administrative console to publish Web Services Description Language (WSDL) files. You can also specify Java Messaging Service (JMS) URL prefixes.

To publish a WSDL file with the administrative console:

1. Open the administrative console.
2. Click **Applications > Enterprise Applications > *application\_instance***. Under Additional Properties, click **Publish WSDL files** which brings you to the **Publish WSDL zip files** panel.
3. **5.1** Specify the default URL prefixes for the Web service.



- a. Select **HTTP URL prefix**.
  - b. Select an entry from the drop down list. If you have multiple application modules, select the application module's check box on the module table.
  - c. Click **Apply**. The URL prefix is copied to the selected module HTTP URL prefix field.
  - d. Click **OK**.
  - e. Click the exported *WSDL\_zip\_file* listed on the **Export WSDL Zip file** panel.
  - f. Download the zip file. Follow your browser's instructions to download the zip file.
4. **5.1** Specify custom URL prefixes for the Web service.
- a. Select **Custom HTTP URL prefix**.
  - b. Type the name of the URL prefix in the **Custom HTTP URL prefix** field. The entry must be of the form `http|https://<host_name>:<port_number>`. For example:  
`http://myHost:999`  
 If you have multiple application modules, select the application module's checkbox on the module table.
  - c. Click **Apply**. The URL prefix is copied to the selected module HTTP URL prefix field.
  - d. Click **OK**.
  - e. Click the exported *WSDL\_zip\_file* listed on the **Export WSDL Zip file** panel.
  - f. Download the zip file. Follow your browser's instructions to download the zip file.
5. **5.1** Specify a JMS URL prefix.
- a. Select the application module.
  - b. Type the JMS URL prefix into the **JMS URL prefix** field. The entry must be of the form:  
`jms:[<queue|topic>?destination=<queue or topic_jndi_name>&connectionFactory=<connection_factory_jndi_name>`. For example:  
`jms:/queue?destination=jms/Q1&connectionFactory=jms/QCF1`
  - c. Click **OK**.
  - d. Click the exported *WSDL\_zip\_file* listed on the **Export WSDL Zip file** panel.
  - e. Download the zip file. Follow your browser's instructions to download the zip file.

*Publish WSDL zip files:*

Use this page to publish Web Services Description Language (WSDL) files.

**5.1** To view this administrative console page, click **Applications >Enterprise Applications > application\_instance > Publish WSDL**.

**5.1** When you click **OK**, a zip file of all the Web services-enabled modules in the application is produced. The name of the published zip file is `application_name_WSDLFiles.zip`. In the published zip file, the directory structure is `application_name/module_name/[META-INF|WEB-INF]/wsdl/wsd_file_name`.

**5.1** In a published WSDL file, the location attribute of a service `soap:address` stanza contains the URL through which the Web service is accessed. You can specify the portion of the URL to be used for the Web services in each module. You can access the Web services in a module through a HTTP transport or JMS transport, or both. You can specify URL information for both types of transports.

*Specify URL prefixes for Web Services:*

Specifies the *protocol* (either `http` or `https`), *host\_name*, and *port\_number* to be used in the URL.

The URL prefix format is `protocol://host_name:port_number`, for example, `http://myHost:9045`. The actual URL that appears in a published WSDL file consists of the prefix prepended to the module's context-root and the Web service url-pattern, for example, `http://myHost:9045/services/myService`.

#### *Select HTTP URL prefix:*

Specifies the drop down list associated with a default list of URL prefixes. This list is the intersection of the set of ports for the module's virtual host and the set of ports for the module's application server. Use items from this list if the Web services application server is accessed directly.

To set an HTTP prefix, select either the **HTTP URL prefix** or **Custom HTTP URL prefix**, enter the value, select the check box of the modules that are to use the prefix, and click **Apply**. When you click **Apply**, the entry in the **Select HTTP URL prefix** or **Custom HTTP URL prefix** fields, depending on which is selected, is copied into the **HTTP URL prefix** field of any module whose check box (in the leftmost column) is selected. The HTTP prefix is not applied to the fields in the JMS URL prefix column.

#### *Custom HTTP URL prefix:*

Specifies the *protocol*, *host*, and *port\_number* of the intermediate service if the Web services in a module are accessed through an intermediate node, for example the Web services gateway or an IHS server.

To set an HTTP prefix, select either the **HTTP URL prefix** or **Custom HTTP URL prefix**, enter the value, select the check box of the modules that are to use the prefix, and click **Apply**. When you click **Apply**, the entry in the **Select HTTP URL prefix** or **Custom HTTP URL prefix** fields, depending on which is selected, is copied into the **HTTP URL prefix** field of any module whose check box (in the leftmost column) is selected. The HTTP prefix is not applied to the fields in the JMS URL prefix column.

#### *JMS URL prefix:*

Specifies the JMS URL prefix string used for each module.

The URL prefix specified must contain the destination and connectionFactory properties. It can contain other property-value pairs, but it must not contain the targetService property, which is added by the system when the published WSDL files are created. The format of the JMS URL prefix is `jms:[queue&topic]?destination=target_queue_or_topic_jndi_name&connectionFactory=factory_jndi_name`, for example,

`jms:/queue?destination=jms/Q1&connectionFactory=jms/QCF`. The actual URL that appears in a published WSDL file consists of the prefix prepended to the Web service targetService, for example, `jms:/queue?destination=jms/Q1&connectionFactory=jms/QCF&targetService=StockQuote`.

#### *Web services client bindings:*

Use this page to specify the Web Service Description Language (WSDL) file name and port type mappings for the Web services in a module.

To view this page, click **Applications >Enterprise Applications > application\_instance > Web Modules > module\_instance>Web Services Client Bindings**.

For EJB modules, click **Applications >Enterprise Applications > application\_instance > EJB Modules > module\_instance>Web Services Client Bindings**

#### *Web Service:*

Specifies the name of this Web service. A module can contain one or more Web services.

*URI:*

Specifies the Uniform Resource Identifier (URI) of the binding file that defines the scope. The URI is relative to the module.

*WSDL Filename:*

Specifies the WSDL file name, which is relative to the module.

A Web service can specify the relative path within the module of a compatible WSDL file containing the actual URL to be used for requests. This is needed only if the original WSDL file does not contain a URL or when a different URL is needed. For a service endpoint with multiple ports defined, a default port mapping specifies the port to use for a port type.

*Default Port Mappings:*

Specifies and manages the default port type mapping for a Web service when a particular port type is requested.

*Web services implementation scope:*

Use this page to view and manage the scope of the ports of a Web Service.

To view this administrative console page, click **Applications >Enterprise Applications > application\_instance > Web Modules > module\_instance>Web Services Implementation Scope**.

*Port:*

Specifies a port name for a Web service. A module can contain one or more Web services, each of which can contain one or more ports.

*Web Service:*

Specifies the name of the Web service. A module can contain one or more Web services.

*URI:*

Specifies the Uniform Resource Identifier (URI) of the binding file that defines the scope. The URI is relative to the Web module.

*Scope:*

Specifies the scope of a port.

The scope determines when a new instance of a service implementation is created for the Web service ports in a module. An application scope causes the same instance of the implementation to be used for all requests on the application. A session scope causes the same instance to be used for all requests on each session. A request scope causes a new instance to be used on every request.

*Publishing WSDL files using the wsadmin command:*

The Web Services Description Language (WSDL) files in each Web services-enabled module are published to the file system location you specify. You can provide these WSDL files to the clients that want to invoke your Web services.

The scripting client, **wsadmin**, can publish the WSDL files in either local, for example, -conntype NONE, or remote mode. However, in local mode, the target application should be located at the same node where the **wsadmin** command is invoked.

The steps below assume that the application has been deployed and that the application server is running.

To publish a WSDL file with the **wsadmin** command:

1. From a command prompt, start `install_root\bin\wsadmin` if you are using Windows or `install_root/bin/wsadmin` if you are using UNIX.
2. At the **wsadmin** command prompt, enter one of the two commands:
  - **\$AdminApp publishWSDL *app\_Name path\_Name***
  - **\$AdminApp publishWSDL *app\_Name path\_Name* {{module {{*binding url-prefix*}}}}**

Where

- *app\_Name* is the application name
- *path\_Name* is the absolute path to the zip file that will contain the published WSDL files. The zip file is saved on the machine running WebSphere Application Server, therefore, if the server is running on a different machine, you need to obtain the zip file from that machine. The directory structure of the resulting zip file is based on the following information:

```
Application file name
 module file name
 META-INF/ or WEB-INF/
 wsdl/
 WSDL file name
```

See the usage scenario for an example of this directory structure.

- *binding* is either http or jms (both are in lower case)
- *url-prefix* is the partial SOAP address for the associated SOAP binding. For an HTTP binding the form is `http://host:port/` or `https://host:port`.

For Java Message Service (JMS) bindings, the form is  
`jms:/queue?destination=dest&connectionFactory=cf` or  
`jms:/topic?destination=dest&connectionFactory=cf`

The **\$AdminApp publishWSDL *app\_Name path\_Name*** command updates the WSDL SOAP address prefixes using the default values. If you do not want to update the WSDL SOAP address prefixes, use the other command, instead of the default values.

The **\$AdminApp publishWSDL *app\_Name path\_Name* {{module {{*binding url-prefix*}}}}** command allows you to customize the WSDL SOAP address for each module. You can specify a different address prefix for each SOAP binding.

The WSDL files from Web services are published to a specified zip file. You can hand the zip file to the client and the client can use the published WSDL files to create a Web services client that accesses the deployed service.

The command to publish WSDL files for a Web service named `WebServicesSamples` could be  
**\$AdminApp publishWSDL WebServicesSamples c:/temp/samplesWSDL.zip**

or

**\$AdminApp publishWSDL WebServicesSamples c:/temp/samplesWSDL.zip { {AddressBookJ2WB.war  
{{http http://localhost:9080}} {StockQuote.jar {{http https://localhost:9443}} } }**

The directory structure for this created zip files is

WebServicesSamples.ear/StockQuote.jar/META-INF/wsdl/StockQuoteFetcher.wsdl  
WebServicesSamples.ear/AddressBookW2JE.jar/META-INF/wsdl/AddressBookW2JE.wsdl  
WebServicesSamples.ear/AddressBookJ2WE.jar/META-INF/wsdl/AddressBookJ2WE.wsdl  
WebServicesSamples.ear/AddressBookJ2WB.war/WEB-INF/wsdl/AddressBookJ2WB.wsdl  
WebServicesSamples.ear/AddressBookW2JB.war/WEB-INF/wsdl/AddressBookW2JB.wsdl

### *Publishing WSDL files using a URL:*

Before you can publish a Web Services Description Language (WSDL) file using a URL, the Web services-enabled application should be installed and running.

The files referenced by the `<wsdl-file>` element in the `webservices.xml` deployment descriptor file can or cannot import other WSDL or XML schema definition (XSD) files. Typically, all WSDL or XSD files are originally placed into the `META-INF/wsdl` directory when using enterprise JavaBeans (EJBs) or the `WEB-INF/wsdl` directory when using Java beans. If your WSDL or XSD files are not placed in one of these directories, the file referenced by the `<wsdl-file>` element and its imported files are located at the same directory and copied to the `wsdl/` directory for publishing purposes.

**Note:** **5.1** EJB-based Web service applications must have an HTTP router or a Web module. Only HTTP URLs are supported for publishing.

To publish a WSDL file using a URL:

1. Retrieve the outer-most WSDL file. The outer-most WSDL file is the WSDL file defined by the `<wsdl-file>` element in the `webservices.xml` deployment descriptor file.  
Each Web service has an endpoint address, for example, `http://example.com/services/stockquote`. You can retrieve the outer-most WSDL file by appending the string with `/wsdl` or `/wsdl/` to the endpoint address, for example, `http://example.com/services/stockquote/wsdl`. **5.1+** You get a WSDL back that supports all protocols, including HTTP, Java Messaging Service and capabilities to directly access EJB bindings.
2. Retrieve the imported WSDL files. When the outer-most WSDL file imports other WSDL or XSD files, these imported files can be retrieved by appending the relative path to the URL, which is used to retrieve the outer-most WSDL file. This is also true for WSDL files that import other files. This process is similar to the use of relative hyperlinks in Hypertext Markup Language (HTML) documents. If an HTML document contains a hyperlink to other documents, the relative path is appended to create the URL to access the hyperlinked documents.

Suppose you have an application with the following directory structure:

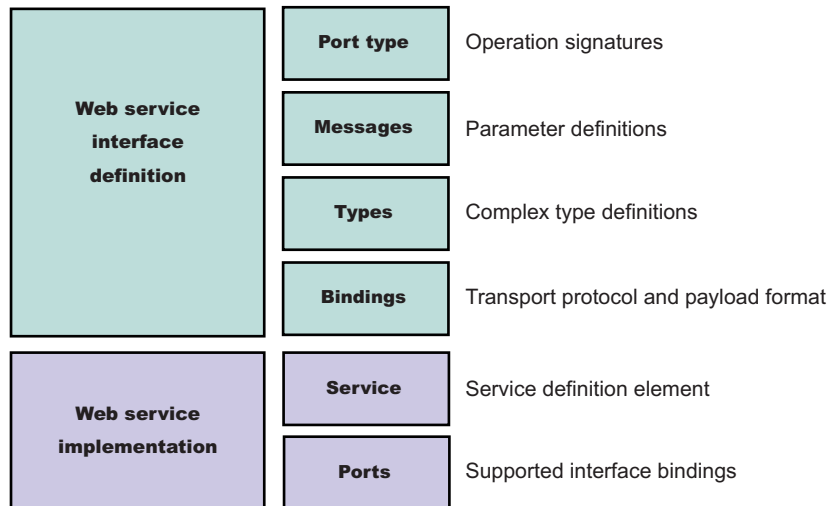
```
<module-root>/
META-INF/
WEB-INF/
webservices.xml/* define Foo service, the <wsdl-file> element points to "/wsdl/fooImpl.wsdl"
*/ web.xml
ibm-webservices-bnd.xml
<jaxrpc-mapping-file>
wsdl/
fooImpl.wsdl/* importing foo.wsdl which is an interface wsdl */
foo.wsdl /* importing type definition for the interface */
fooTypes.xsd
```

If the SOAP address for the foo service is `http://examples.com:9080/services/foo`, the simple way to retrieve the foo service's outer-most WSDL is with the following form:

`http://examples.com:9090/services/foo/wsdl` or `http://examples.com:9090/services/foo/wsdl/`. The URL is redirected to `http://examples.com:9090/services/foo/wsdl/fooImpl.wsdl`, where `fooImpl.wsdl` is the name of the outer-most WSDL file.

Since the `fooImpl.wsdl` file has the import `<import namespace="http://examples.com/foo" location="a/b/foo.wsdl">`, use the URL `http://examples.com:9090/services/foo/wsdl/a/b/foo.wsdl` to obtain the `foo.wsdl` file.

**WSDL anatomy:** Web Services Description Language (WSDL) files are written in eXtensible Markup Language (XML). To learn more about XML, see [Web services: Resources for learning](#).



A WSDL contains the following parts:

- **Web service interface definition**  
This is where the elements are contained, as well as the namespaces.
- **Web service implementation**  
This is where you find the definition of the service and ports.

A WSDL file describes a Web service with the following elements:

### portType

The description of the operations and their associated messages. PortTypes define abstract operations.

```
<portType name="EightBall">
 <operation name="getAnswer">
 <input message="ebs:IngetAnswerRequest"/>
 <output message="ebs:OutgetAnswerResponse"/>
 </operation>
</portType>
```

### message

The description of parameters (input and output) and return values.

```
<message name="IngetAnswerRequest">
 <part name="meth1_inType" type="ebs:questionType"/>
</message>
<message name="OutgetAnswerResponse">
 <part name="meth1_outType" type="ebs:answerType"/>
</message>
```

## types

The schema for describing XML complex types used in the messages.

```
<types>
 <xsd:schema targetNamespace="...">
 <xsd:complexType name="questionType">
 <xsd:element name="question" type="string"/>
 </xsd:complexType>
 <xsd:complexType name="answerType">
 ...
 </xsd:complexType>
 </types>
```

## binding

Bindings describe the protocol used to access a service, as well as the data formats for the messages defined by a particular portType.

```
<binding name="EightBallBinding" type="ebs:EightBall">
 <soap:binding style="rpc" transport="schemas.xmlsoap.org/soap/http">
 <operation name="ebs:getAnswer">
 <soap:operation soapAction="urn:EightBall"/>
 <input>
 <soap:body namespace="urn:EightBall" ... />
 </input>
 ...
</binding>
```

The remaining parts, services and ports, indicate where you can find the WSDL.

## Service

Contains the Web service name and a list of the ports.

## Ports

Contains the location of the Web service and the binding to be used to access the service.

```
<service name="EightBall">
 <port binding="ebs:EightBallBinding" name="EightBallPort">
 <soap:address location="localhost:8080/axis/EightBall"/>
 </port>
</service>
```

**Multipart WSDL best practices:** WebSphere Application Server supports deployment of Web services using a multipart Web Services Description Language (WSDL) file. That is, WSDL files import other WSDL files when the WSDL file listed in the `<wsdl-file>` element of the `webservices.xml` deployment descriptor contains all `<wsdl:service>` and `<wsdl:port>` elements. The WSDL file is divided into an implementation WSDL and an interface WSDL.

The `<wsdl:import>` element indicates a reference to another WSDL file. If the `<wsdl:import>` element location attribute does not contain a URL, that is, it contains only a file name, and does not begin with `http://`, `https://` or `file://`, the imported file must be located in the same directory and must not contain a relative path component. For example, if `META-INF/wsdl/A_Impl.wsdl` is in your module and contains the import statement `<wsdl:import="A.wsdl" namespace="...">`, the file, `A.wsdl` must also be located in the module `META-INF/wsdl` directory.

It is recommended that all WSDL files be placed in either the `META-INF/wsdl` directory, if you are using enterprise JavaBeans (EJBs), or the `WEB-INF/wsdl` directory, if you are using Java beans, even if there are relative imports within the WSDL files. Otherwise, there are implications with the WSDL publication when you use a path like the following `<location=" ../interfaces/A_Interface.wsdl" namespace="...">`. Using a path like this fails because the presence of the relative path, regardless of whether the file is located at that path or not. If the location is a URL, it must be readable at both deployment and server startup.



## WSDL publication

The files located in the META-INF/wsdl or WEB-INF/wsdl directory can be published through either a URL or file, including WSDL or XSD files. For example, if the file referenced in the `<wsdl:file>` element of the `webservices.xml` deployment descriptor is located in the META-INF/wsdl or WEB-INF/wsdl directory, it is publishable. If the files imported by the `<wsdl:file>` are located in the `wsdl/` directory or its subdirectory, they are publishable.

If the WSDL file referenced by the `<wsdl:file>` element is located in a directory other than `wsdl/`, or its subdirectories, the file and its imported files, either WSDL or XSD files, which are in the same directory, are copied to the `wsdl` directory without modification when the application is installed. These types of files can also be published.

If the `<wsdl:file>` imports a file located in a different directory, the file is not copied to the `wsdl` directory and not available for publishing.

## Developing a Service Endpoint Interface for a Java bean implementation

Access an existing Java bean WAR file.

Set up a Web services development and unmanaged client execution environment.

The Service Endpoint Interface defines the methods for a particular Web service. The Java bean implementation must implement methods having the same signature as the methods on the Service Endpoint Interface. There are a number of restrictions on which types to use as parameters and results of Service Endpoint Interface methods. These restrictions are documented in the Java API for XML-based remote procedure call (JAX-RPC) specification, which is available through Web services: Resources for learning.

You can also create a Service Endpoint Interface by using the Assembly Toolkit, which is a component of the Application Assembly Toolkit. The steps are similar except the Assembly Toolkit automatically compiles the interface when you save it.

To develop a Service Endpoint Interface for a Java bean implementation:

1. Create a Java interface containing the methods to include in the Service Endpoint Interface. The interface should extend the `java.rmi.Remote` interface. Each method throws the exception, `java.rmi.RemoteException`. If you start with an existing Java interface, remove any methods that do not conform to JAX-RPC. See the example that follows this tasks steps.
2. Compile the interface. Use the **javac** commands for Windows and UNIX platforms listed in the topic Developing thin application client code to compile the interface. In the **javac** command, use the name of the Service Endpoint Interface class for the class to be compiled.

A Service Endpoint Interface which you can use to develop a Web service.

This example uses a Java interface called `AddressBook`. The following example depicts the `AddressBook` interface:

```
package addr;
public interface AddressBook {
 /**
 * Retrieve an entry from the AddressBook.
 *
 * @param name the name of the entry to look up.
 * @return the AddressBook entry matching name or null if none.
 * @throws java.rmi.RemoteException if communications failure.
 */
 public addr.Address getAddressFromName(java.lang.String name);
}
```

You use the AddressBook Java interface to create the Service Endpoint Interface:

1. Begin with the remote interface, AddressBook.java.
2. Make a copy of the remote interface named AddressBook\_SEI.java and use it as a template for the Service Endpoint Interface.
3. Change the interface to extend the java.rmi.Remote interface.
4. Modify each method declaration to add a throws clause for java.rmi.RemoteException.
5. Compile the interface.

Use the Service Endpoint Interface to Develop a Web Services Description Language (WSDL) file.

## Developing Web services deployment descriptor templates for a Java bean implementation

To develop the deployment descriptor templates from a Web Services Description Language (WSDL) file, you must obtain the Uniform Resource Locator (URL) of the WSDL file to use.

If it is a local file and you are running the Windows platform, the URL looks like this:

*file:drive:\path\file\_name.wsdl*. If you are using the UNIX platform, the URL looks like this:

*file:/path/file\_name.wsdl*. You can also specify local files using the absolute or relative file system path.

When the Web service implementation is a Java bean in a Web module, the webservices.xml, ibm-webservices-bnd.xmi and ibm-webservices-ext.xmi deployment descriptors and the Java API for XML-based remote procedure call (JAX-RPC) mapping file are generated in the WEB-INF subdirectory.

To develop deployment descriptor templates:

Run the **WSDL2Java -verbose -role develop-server -container web -genJava no wsdlURL** command to generate the server deployment descriptor templates and mapping file into the WEB-INF subdirectory. If the **-verbose** option is specified, a list of all generated files displays when the command runs.

Deployment descriptor templates that are required to implement or use a Web service.

The following example uses a WSDL file named AddressBookJ2WB.wsdl:

1. Generate the template files:
  - `WSDL2Java -verbose -role develop-server -container web -genJava no AddressBookJ2WB.wsdl`

The deployment descriptor templates and mapping file are generated into the WEB-INF subdirectory as follows:

```
Parsing XML file: AddressBookJ2WB.wsdl
Generating: WEB-INF\webservices.xml
Generating: WEB-INF\ibm-webservices-bnd.xmi
Generating: WEB-INF\ibm-webservices-ext.xmi
Generating: WEB-INF\AddressBookJ2WB_mapping.xml
```

## Developing a Web service using a stateless session enterprise bean

Set up a Web services development and unmanaged client execution environment.

To use an enterprise bean as the basis for a Web service implementation, follow these requirements:

- The enterprise bean must be a stateless session bean.
- Web service method parameters must be serializable and cannot be object references.
- Web service method parameters must be one of the supported Java API for XML-based remote procedure call (JAX-RPC) types.

These requirements are documented in the JAX-RPC specification available through Web services: Resources for learning.

Create the artifacts that enable the enterprise bean to be a Web service and assemble the artifacts into the enterprise application as follows:

1. Access an existing Java archive (JAR) file to be used as a Web service. Make sure that the enterprise bean meets the requirements.
2. Develop an EJB Service Endpoint Interface. The Service Endpoint Interface defines which enterprise bean methods should be made available as a Web service.
3. Develop a Web Services Description Language (WSDL) file.
4. Develop Web services deployment descriptor templates from an EJB implementation.
5. Assemble a Web services-enabled JAR file.
6. Configure the `webservices.xml` deployment descriptor.
7. Configure the `ibm-webservices-bnd.xmi` deployment descriptor.
8. Assemble a Web services-enabled enterprise archive (EAR) file.
9. Enable the EAR file. When the EAR file contains EJB modules, it must have the Web services endpoint Web archive (WAR) file added with the `endptEnabler` tool before it is deployed.
10. Deploy the EAR file into WebSphere Application Server.

A Web service from a stateless session enterprise bean.

## Developing a Service Endpoint Interface from an EJB remote interface

Set up a Web services development and unmanaged client execution environment.

The Service Endpoint Interface defines the Web services methods. The enterprise JavaBean (EJB) that implements the Web service must implement methods having the same signature as the methods of the Service Endpoint Interface. There are a number of restrictions on which types to use as parameters and results of Service Endpoint Interface methods. These restrictions are documented in the Java API for XML-based remote procedure call (JAX-RPC) specification, which is available through Web services: Resources for learning.

The easiest method for creating the Service Endpoint Interface for an EJB Web service implementation is from the EJB remote interface.

You can also create a Service Endpoint Interface by using the Assembly Toolkit, which is a component of the Application Server Toolkit. The steps are similar except the Assembly Toolkit automatically compiles the interface when you save it.

To develop a Service Endpoint Interface:

1. Create a Java interface containing the methods to include in the Service Endpoint Interface. The interface should extend the `java.rmi.Remote` interface. Each method throws the exception, `java.rmi.RemoteException`. If you start with an existing Java interface, remove any methods that do not conform to JAX-RPC.
2. Compile the interface. Use the **javac** commands for Windows and UNIX platforms listed in the topic Developing thin application client code to compile the interface. In the **javac** command, use the name of the Service Endpoint Interface class for the class to be compiled.

A Service Endpoint Interface which you can use to develop a Web service.

This example uses an EJB remote interface called `AddressBook_RI`.

```
package addr;
public interface AddressBook_RI extends javax.ejb.EJBObject {
 /**
 * Retrieve an entry from the AddressBook.
 *
 * @param name the name of the entry to look up.
 */
}
```

```

 *@return the AddressBook entry matching name or null if none.
 *@throws java.rmi.RemoteException if communications failure.
 */
 public addr.Address getAddressFromName(java.lang.String name)
 throws java.rmi.RemoteException;
}

```

You use the AddressBook\_RI remote interface to create the Service Endpoint Interface:

1. Begin with the remote interface, AddressBook\_RI.java:
2. Make a copy of the remote interface named AddressBook.java and use it as a template for the Service Endpoint Interface.
3. Change the interface to extend the java.rmi.Remote interface, instead of the javax.ejb.EJBObject Service Endpoint Interface.
4. Compile the AddressBook.java Service Endpoint Interface.

Use the Service Endpoint Interface to Develop a WSDL file.

## Developing Web services deployment descriptor templates for an EJB implementation

To develop the deployment descriptor templates from a Web Services Description Language (WSDL) file, you must obtain the Uniform Resource Locator (URL) of the WSDL file to use.

If it is a local file and you are running the Windows platform, the URL looks like this:  
*file:drive:\path\file\_name.wsdl*. If you are using the UNIX platform, the URL looks like this:  
*file:/path/file\_name.wsdl*. You can also specify local files using the absolute or relative file system path.

When the Web service implementation is an enterprise Java bean (EJB) in an EJB module, the webservices.xml, ibm-webservices-bnd.xmi and ibm-webservices-ext.xmi deployment descriptors, and the Java API for XML-based remote procedure call (JAX-RPC) mapping file are generated in the META-INF subdirectory.

To develop deployment descriptor templates:

Run the **WSDL2Java -verbose -role develop-server -container ejb -genJava no wsdIURL** command to generate the server deployment descriptor templates and mapping file into the META-INF subdirectory. If the **-verbose** option is specified, a list of all generated files displays when the command runs.

Deployment descriptor templates that are required to implement a Web service.

The following example uses a WSDL file named AddressBookJ2WE.wsdl:

1. Generate the template files:
  - WSDL2Java -verbose -role develop-server -container ejb -genJava no AddressBookJ2WE.wsdl

The deployment descriptor templates are generated into the META-INF subdirectory as follows:

```

Parsing XML file: AddressBookJ2WE.wsdl
Generating: META-INF\webservices.xml
Generating: META-INF\ibm-webservices-bnd.xmi
Generating: META-INF\ibm-webservices-ext.xmi
Generating: META-INF\AddressBookJ2WE_mapping.xml

```

## Completing the EJB implementation

Develop EJB implementation templates and bindings from a Web Services Description (WSDL) file.

To complete the EJB implementation:

1. Inspect the enterprise EJB remote interface template, *portType\_RI.java*. If necessary, modify the template. *portType* is the name of the <wsdl:portType> element in the WSDL file.
2. Inspect the EJB home interface template, *portTypeHome.java*. If necessary, modify the template.

3. Edit the EJB implementation template, *bindingImpl.java*. *binding* is the name of the <wsdl:binding> element in the WSDL file.
  - a. Complete the implementation of the methods in the template.
  - b. (Optional) Make changes if necessary.
  - c. (Optional) Change the class name if the binding name is not acceptable.
4. Compile all the Java classes.
5. Assemble an EJB Java archive (JAR) file. Assemble all the Java classes into an EJB JAR file using the typical EJB assembly tools. Include all of the classes generated from running the **WSDL2Java** command tool when developing implementation templates and bindings from a WSDL file.

An EJB JAR file containing an EJB and supporting classes created from a WSDL file.

Configure the webservices.xml deployment descriptor .

## Developing a new Web service with an existing WSDL file using a Java bean

Locate the Web Services Description Language (WSDL) file that defines the Web service to be implemented. You can develop a WSDL or obtain one from an existing Web service through e-mail, downloading or a Uniform Resource Locator (URL).

To develop a new Web service with an existing WSDL file using a Java bean:

1. Develop Java bean implementation templates and bindings from a WSDL file.
2. Complete the Java bean implementation.
3. Assemble a Web services-enabled Web archive (WAR) file when starting from a WSDL file.
4. Configure the webservices.xml deployment descriptor.
5. Configure the ibm-webservices-bnd.xmi deployment descriptor.
6. Assemble a Web services-enabled WAR into an EAR file.
7. Deploy the EAR file into WebSphere Application Server.

Develop Web services deployment descriptor templates from a WSDL file.

You can either develop Web services deployment descriptor templates for a Java bean implementation or develop Web services deployment descriptor templates for an EJB implementation.

## Developing Web services deployment descriptor templates for a Java bean implementation

To develop the Java bean implementation templates and bindings from a Web Services Description (WSDL) file, you must obtain the Uniform Resource Locator (URL) of the WSDL file to use.

If it is a local file and you are running the Windows platform, the URL looks like this: `file:drive:\path\file_name.wsdl`. If you are using the UNIX platform, the URL looks like this: `file:/path/file_name.wsdl`. You can also specify local files using the absolute or relative file system path.

Implementation templates are generated using the `-role develop-server` option of the **WSDL2Java** command. The **WSDL2Java** command also generates bindings and deployment descriptors.

To develop Java bean implementation templates and bindings from a WSDL file: Run the **WSDL2Java -verbose -role develop-server -container web wsdIURL** command. Since the **verbose** option is specified, a list of all generated files is displayed when the command runs.

Templates for the implementation and deployment descriptors required to implement a Web service, as well as bindings files. These templates are partially filled with information from the WSDL file.

The following example uses an Java bean named `AddressBook` and a WSDL file named `AddressBook.wsdl`. After generating the template files from the **WSDL2Java -verbose -role develop-server -container web AddressBook.wsdl** command, the following files are generated:

```
Parsing XML file: file:e:/example/app/topdown/step1/AddressBook.wsdl
WSWS3185I: Info: Parsing XML file: AddressBook.wsdl
WSWS3282I: Info: Generating addr\Address.java.
WSWS3282I: Info: Generating addr\Phone.java.
WSWS3282I: Info: Generating addr\StateType.java.
WSWS3282I: Info: Generating addr\AddressBook.java.
WSWS3282I: Info: Generating addr\AddressBookSoapBindingImpl.java..
WSWS3282I: Info: Generating WEB-INF\webservices.xml.
WSWS3282I: Info: Generating WEB-INF\ibm-webservices-bnd.xmi.
WSWS3282I: Info: Generating WEB-INF\AddressBook_mapping.xml.
WSWS3282I: Info: Generating WEB-INF\ibm-webservices-ext.xmi.
```

The generated file named `AddressBookSOAPBindingImpl.java` is the template for the implementation bean. It is named after the port in the WSDL file. Generally, this class is renamed to a more meaningful name.

Complete the Java bean implementation.

## Completing the Java bean implementation

Develop Java bean implementation templates and bindings from a Web Services Description Language (WSDL) file.

1. Edit the Java bean implementation template, `bindingImpl.java`. *binding* is the name of the `<wsdl:binding>` element in the WSDL file.
  - a. Complete the implementation of the methods in the template.
  - b. (Optional) Make changes if necessary.
  - c. (Optional) Change the class name if the binding name is not acceptable.
2. Compile all the Java classes.
3. Assemble a Web archive (WAR) file. Assemble all the Java classes into a WAR file using typical Web module assembly tools. Include all of the classes generated from running the **WSDL2Java** command tool when developing implementation templates and bindings from a WSDL file.

A Java archive (JAR) file containing a Java bean and supported classes created from the WSDL file.

Configure the `webservices.xml` deployment descriptor.

## Developing a new Web service from an existing WSDL file using a stateless session enterprise bean

Set up a Web services development and unmanaged client execution environment.

Locate the Web Services Description Language (WSDL) file that defines the Web service to implement. The SOAP address URI is not required because it is updated when your new implementation is deployed.

Create the enterprise bean and artifacts that enable the enterprise bean to be a Web service and assemble those artifacts into the enterprise application as follows:

1. Develop implementation templates and bindings from a WSDL file.
2. Complete the enterprise bean implementation.
3. Assemble a Web services-enabled enterprise EJB Java archive (JAR) file.
4. Configure the `webservices.xml` deployment descriptor.



5. Configure the `ibm-webservices-bnd.xmi` deployment descriptor.
6. Assemble a Web services-enabled EJB JAR into an EAR file.
7. Enable the EAR file. When the EAR file contains EJB modules, it must have the Web services endpoint Web archive (WAR) file added with the **endptEnabler** command or Assembly Toolkit before deployment.
8. Deploy the EAR file into WebSphere Application Server.

An EJB implementation of a Web service defined in the WSDL file.

## Developing EJB implementation templates and bindings from a WSDL file

To develop enterprise JavaBean (EJB) implementation templates and bindings from a Web Services Description (WSDL) file, you must obtain the Uniform Resource Locator (URL) of the WSDL file to use.

If it is a local file and you are running the Windows platform, the URL looks like this: `file:drive:\path\file_name.wsdl`. If you are using the UNIX platform, the URL looks like this: `file:/path/file_name.wsdl`. You can also specify local files using the absolute or relative file system path.

Implementation templates are generated using the `-role develop-server` option of the **WSDL2Java** command.

Templates are generated for an EJB implementation for the following:

- EJB
- EJB remote interface
- EJB Home

The **WSDL2Java** command also generates bindings and deployment descriptors.

To develop implementation templates and bindings from a WSDL file: Run the **WSDL2Java -verbose -role develop-server -container ejb wsdIURL** command. Since the **verbose** option is specified, a list of all generated files is displayed when the command runs.

Templates for the implementation and deployment descriptors required to implement a Web service, as well as bindings files. These templates are partially filled with information from the WSDL file.

The following example uses an enterprise bean named `AddressBook` and a WSDL file named `AddressBook.wsdl`. After generating the template files from the **WSDL2Java -verbose -role develop-server -container EJB AddressBook.wsdl** command, the following files are generated:

```
Parsing XML file: file:e:/example/app/topdown/step1/AddressBook.wsdl
WSWS3185I: Info: Parsing XML file: AddressBook.wsdl
WSWS3282I: Info: Generating addr\Address.java.
WSWS3282I: Info: Generating addr\Phone.java.
WSWS3282I: Info: Generating addr\StateType.java.
WSWS3282I: Info: Generating addr\AddressBook.java.
WSWS3282I: Info: Generating addr\AddressBookSoapBindingImpl.java.
WSWS3282I: Info: Generating addr\AddressBook_RI.java.
WSWS3282I: Info: Generating addr\AddressBookHome.java.
WSWS3282I: Info: Generating META-INF\webservices.xml.
WSWS3282I: Info: Generating META-INF\ibm-webservices-bnd.xmi.
WSWS3282I: Info: Generating META-INF\AddressBook_mapping.xml.
WSWS3282I: Info: Generating META-INF\ibm-webservices-ext.xmi.
```

Complete the EJB implementation.



---

## Developing Web services clients based on Web Services for J2EE

This topic explains how to develop a Web services client based on the Web Services for Java 2 platform, Enterprise Edition (J2EE) specification.

For a Java application to act as Web service client, a mapping between the Web Services Description Language (WSDL) file and the Java application must exist. The mapping is defined by the Java API for XML-based RPC (JAX-RPC) specification. You can use a Java component to implement a Web service by specifying the component's interface and binding information in the WSDL file and designing the application server infrastructure to accept the service request. This entire process encompassed is based on the Web Services for J2EE specification.

Before you begin this task, locate the Web Services Description Language (WSDL) file that defines the Web service to access.

To create the client code and artifacts that enable the application client to access a Web service:

1. Develop client bindings from a WSDL file. The client-side bindings and deployment descriptors are generated.
2. Complete the client implementation.
3. (Optional) Assemble a Web services-enabled client Java archive (JAR) file. Complete this step if you are developing a managed client that runs in the J2EE client container.
4. (Optional) Assemble a Web services-enabled client Web archive (WAR) file. Complete this step if you are developing a managed client that runs in the J2EE client container.
5. (Optional) Configure the `webservicesclient.xml` deployment descriptor. Complete this step if you are developing a managed client that runs in the J2EE client container.
6. (Optional) Configure the `ibm-webservicesclient-bnd.xmi` deployment descriptor. Complete this step if you are deploying a managed client that runs in the J2EE client container and you want to override the default client settings. See `ibm-webservicesclient-bnd.xmi` assembly properties for more information about the `ibm-webservicesclient-bnd.xmi` deployment descriptor.
7. Test the Web services-enabled client application.

You have created and tested a Web services client application. For step-by-step information see Example: Developing Web services clients based on Web Services for J2EE.

### Example: Developing Web services clients based on Web Services for J2EE

This example takes you through the steps to develop a Web service client. The development process is based on the Web Services for Java 2 platform, Enterprise Edition (J2EE) and the Java API for XML-based remote procedure call (JAX-RPC) specification. For a Java or J2EE application to act as a client of a Web service, you must map the WSDL file to the Java code. The JAX-RPC specification defines the mapping between a WSDL file, Java code and XML Schema types.

#### Steps for this example task

1. Obtain the Web Services Description Language (WSDL) document for the Web service that you want to access.

You can obtain the WSDL document from the service provider by e-mail or by looking it up in a Universal Description, Discovery and Integration (UDDI) registry.

2. Develop client bindings from your WSDL file.

The WSDL document is used to generate all the information needed to invoke the Web service, including the Service Endpoint Interface and implementations; generated service interface; `webservicesclient.xml` and `ibm-webservicesclient-bnd.xmi` and `ibm-webservicesclient-ext.xmi` deployment descriptors.

The **WSDL2Java** command-line tool is run against your WSDL file to develop client bindings.

3. Implement the client.

See Chapter 4 of the JSR-109 specification. You can access the specification through Web services: Resources for learning.

You can also review the GetQuote sample available in the Samples Gallery. (See "Samples Gallery" in the Information Center.)

4. Assemble the module.

Assemble the client JAR file into an EAR file or assemble the client WAR file into an EAR file.

5. Configure the deployment descriptors.

Configure the `webservicesclient.xml` deployment descriptor.

Configure the `ibm-webservicesclient-bnd.xmi` deployment descriptor.

6.

7. Test the Web services client.

You should test the client to make sure it correctly operates and binds to the Web service.

## Developing client bindings from a WSDL file

To develop the client bindings from a Web Services Description (WSDL) file, you must obtain the Uniform Resource Locator (URL) of the WSDL file to use.

If it is a local file and you are running the Windows platform, the URL looks like this:

`file:drive:\path\file_name.wsdl`. If you are using the UNIX platform, the URL looks like this:

`file:/path/file_name.wsdl`. You can also specify local files using the absolute or relative file system path.

Client bindings are generated using the `-role develop-client` option in combination with the `-container` option of the `WSDL2Java` command. The `-container` option takes the following parameters:

- **-container client**

Generates bindings and deployment descriptors for a client residing the application client container.

- **-container ejb**

Generates bindings and deployment descriptors for a client that is an EJB in the EJB module.

- **-container web**

Generates bindings and deployment descriptors for a client residing in the Web container.

To develop client bindings from a WSDL file:

Run the **WSDL2Java -verbose -role develop-client -container *type* *wsdlURL*** command.

Where *type* is **ejb** for an enterprise JavaBean (EJB) client, **web** for a Java bean client, or **client** for an application client.

**Note:** You can have:

- `-container web`
- `-container ejb`
- `-container client`

Since the **verbose** option is specified, a list of all generated files is displayed when the command runs.

The bindings and deployment descriptors needed by a client to use a Web service.

The following example uses an enterprise bean named `AddressBook` and a WSDL file named `AddressBook.wsdl`. After generating the bindings from the **WSDL2Java -verbose -role develop-client -container client AddressBook.wsdl** command, the following files are generated:

```
Parsing XML file: file:e:/example/app/topdown/step1/AddressBook.wsdl
WSWS3185I: Info: Parsing XML file: AddressBook.wsdl
WSWS3282I: Info: Generating addr\Address.java.
WSWS3282I: Info: Generating addr\Phone.java.
WSWS3282I: Info: Generating addr\StateType.java.
WSWS3282I: Info: Generating addr\AddressBook.java.
WSWS3282I: Info: Generating addr\AddressBookService.java.
WSWS3282I: Info: Generating META-INF\webservicesclient.xml.
WSWS3282I: Info: Generating META-INF\ibm-webservicesclient-bnd.xmi.
WSWS3282I: Info: Generating META-INF\AddressBook_mapping.xml.
WSWS3282I: Info: Generating META-INF\ibm-webservicesclient-ext.xmi.
```

Complete the client implementation.

Assemble a Web services-enabled client JAR and EAR file.

## Assembling a Web services-enabled client JAR file into an EAR file

You need the following artifacts:

- Assembled client module, containing the implementation, all classes generated by the **WSDL2Java** command-line tool, MANIFEST.MF and deployment descriptor. This module can be:
  - An application client module containing META-INF/application-client.xml
  - An enterprise JavaBean (EJB) module containing META-INF/ejb-jar.xml
- Web Services Description Language (WSDL) file used to develop the client
- Templates for webservicesclient.xml and ibm-webservicesclient-ext.xmi deployment descriptors, if used.
- Generated JAX-RPC mapping deployment descriptor

You can use the Assembly Toolkit to assemble Web service-enabled client applications.

To assemble the client code and artifacts that enable the application client to access a Web service:

1. Start the Assembly Toolkit.
2. Click **File > Import** to import the EJB JAR file, App Client JAR file, or WAR file into the Assembly Toolkit.
3. Open the J2EE perspective by clicking **Windows > Open Perspective > Other > J2EE**.
4. Switch to the **Project Navigator** pane by clicking the **Project Navigator** tab.
5. Locate the project for the file you just imported in the **Project Navigator** pane.
6. Expand the `ejbModule` (for an EJB JAR file) or the `appClientModule` (for the application client JAR file) entry so the META-INF directory is displayed. Expand the META-INF directory.
7. Right-click the META-INF directory and select **New > Folder**. Create a subfolder named `wsdl` in the META-INF directory.
  - a. Copy the WSDL file to the META-INF\wsdl directory by right-clicking on the `wsdl` directory and click **File > Import > File system**. Browse the WSDL file for this Web service and click **Finish**.
  - b. Copy the `webservicesclient.xml` and the JAX-RPC mapping file in the META-INF subdirectory in the same manner you copied the WSDL file. The JAX-RPC mapping file is indicated by the `<jaxrpc-mapping-file>` element in the `webservicesclient.xml` file.
  - c. (Optional) Place the `ibm-webservicesclient-ext.xmi` and the `ibm-webservicesclient-bnd.xmi` file in the META-INF subdirectory, if used.
8. Assemble the JAR file into an EAR file using typical assembly techniques if the client runs in a container.
9. Right-click on the WEB-INF directory and select **New > Folder**. Create a subfolder named `wsdl` in the WEB-INF directory.
  - a. Copy the WSDL file to the WEB-INF\wsdl directory by right-clicking on the `wsdl` directory and click **File > Import > File system**. Browse the WSDL file for this Web service and click **Finish**.

- b. Copy the `webservicesclient.xml` and the JAX-RPC mapping file in the WEB-INF subdirectory in the same manner you copied the WSDL file. The JAX-RPC mapping file is indicated by the `<jaxrpc-mapping-file>` element in the `webservicesclient.xml` file.
- c. (Optional) Place the `ibm-webservicesclient-ext.xmi` and `ibm-webservicesclient-bnd.xmi` file in the WEB-INF subdirectory, if used.

The artifacts required to enable the client module to use Web services are added to the module.

This example uses a JAR file named `AddressBookClient.jar` and an EAR file named `AddressBookClient.ear`:

```

META-INF/MANIFEST.MF
META-INF/application-client.xml
META-INF/wsdl/AddressBook.wsdl
META-INF/webservicesclient.xml
META-INF/AddressBook_mapping.xml
com/ibm/websphere/samples/webservices/addr/Address.class
com/ibm/websphere/samples/webservices/addr/AddressBook.class
com/ibm/websphere/samples/webservices/addr/AddressBookClient.class
com/ibm/websphere/samples/webservices/addr/AddressBookService.class
...other generated classes...

```

After assembling the `AddressBookClient.jar` file into the `AddressBookClient.ear` file, the `AddressBookClient.ear` file contains the following files:

```

META-INF/MANIFEST.MF
AddressBookClient.jar
META-INF/application.xml

```

Configure the `webservicesclient.xml` deployment descriptor .

## Assembling a Web services-enabled client WAR file into an EAR file

You need the following artifacts:

- Assembled client Web archive (WAR) module, containing the implementation, all classes generated by the **WSDL2Java** command-line tool, `MANIFEST.MF` and deployment descriptor.
- Web Services Description Language (WSDL) file used to develop the client
- Templates for `webservicesclient.xml`, `ibm-webservicesclient-bnd.xmi` and `ibm-webservicesclient-ext.xmi` deployment descriptors, if used.
- Generated JAX-RPC mapping deployment descriptor

You can use the Assembly Toolkit to assemble Web service-enabled client applications.

To assemble the client code and artifacts that enable the application client to access a Web service:

1. Start the Assembly Toolkit.
2. Click **File > Import** to import the WAR file into the Assembly Toolkit.
3. Open the J2EE perspective by clicking **Windows > Open Perspective > Other > J2EE**.
4. Switch to the **Project Navigator** pane by clicking the **Project Navigator** tab.
5. Locate the project for the file you just imported in the **Project Navigator** pane.
6. Expand the `webContent` entry so the WEB-INF directory is displayed. Expand the WEB-INF directory.
7. Right-click on the WEB-INF directory and select **New > Folder**. Create a subfolder named `wsdl` in the WEB-INF directory.
  - a. Copy the WSDL file to the `WEB-INF\wsdl` directory by right-clicking on the `wsdl` directory and click **File > Import > File system**. Browse the WSDL file for this Web service and click **Finish**.
  - b. Copy the `webservicesclient.xml` and the JAX-RPC mapping file in the WEB-INF subdirectory in the same manner you copied the WSDL file. The JAX-RPC mapping file is indicated by the `<jaxrpc-mapping-file>` element in the `webservicesclient.xml` file.

- c. (Optional) Place the `ibm-webservicesclient-ext.xml` and `ibm-webservicesclient-bnd.xml` file in the `WEB-INF` subdirectory, if used.
8. Assemble the WAR file into an EAR file using typical assembly techniques.

The artifacts required to enable the client module to use Web services are added to the module.

This example uses a WAR file named `AddressBookWeb.war` and an EAR file named `AddressBook.ear`:

```
WEB-INF/MANIFEST.MF
WEB-INF/web.xml
WEB-INF/wsd1/AddressBook.wsdl
WEB-INF/webservicesclient.xml
WEB-INF/AddressBook_mapping.xml
WEB-INF/ibm-webservicesclient-ext.xml (optional)
WEB-INF/ibm-webservicesclient-bnd.xml
com/ibm/websphere/samples/webservices/addr/Address.class
com/ibm/websphere/samples/webservices/addr/AddressBook.class
com/ibm/websphere/samples/webservices/addr/AddressBookClient.class
com/ibm/websphere/samples/webservices/addr/AddressBookService.class
...other generated classes...
```

After assembling the `AddressBookWeb.war` file into the `AddressBook.ear` file, the `AddressBook.ear` file contains the following files:

```
WEB-INF/MANIFEST.MF
AddressBookWeb.war
WEB-INF/application.xml
```

Configure the `webservicesclient.xml` deployment descriptor .

## Configuring the `ibm-webservicesclient-bnd.xml` deployment descriptor

This topic explains how to configure the `ibm-webservicesclient-bnd.xml` deployment descriptor file using the Assembly Toolkit, which replaces the Application Assembly Tool (AAT). It is one of the tools available with the Application Server Toolkit product.

To configure the `ibm-webservicesclient-bnd.xml` deployment descriptor file:

1. Start the Assembly Toolkit
2. Locate the `webservicesclient.xml` file in the module.
3. Double-click the `webservices.xml` file to open the **Web Services Client** editor.
4. Access the **Web Services Client Bindings** editor through the **Client Binding** tab at the bottom of the editor pane.
5. Verify the `componentNameLink` element settings.
  - a. Open the **Web Services Client Bindings** editor.
  - b. Expand the **Component scoped references** section.
  - c. Click **Add**.
  - d. Select the component scoped references defined in the `webservicesclient.xml` file from the list.
6. Verify the `serviceRefLink` element settings.
  - a. Open the **Web Services Client Bindings** editor.
  - b. Click the **Services References** tab.
  - c. Click **Add**.
  - d. Select the service references defined in the `webservicesclient.xml` file from the list.
7. Verify the `deployWSDLFile` element settings.
  - a. Open the **Web Services Client Bindings** editor.
  - b. Select the service references or component scoped reference desired.

- c. Expand the **Service reference details** section.
- d. Click **Browse** on the **Deployed WSDL file** field.
- e. Select the new WSDL file.
- f. Click **OK**.

The `deployedWSDLFile` element of a deployed Web service can also be changed using the administrative console. Using application management, navigate to the Web module or EJB module of the Web service application and select **Web Services Client Bindings**.

8. Verify the `defaultMappings` element settings.
  - a. Open the **Web Services Client Bindings** editor.
  - b. Click **Default Mappings**.
  - c. Click **Add**.
  - d. Edit the entries in the newly added row to establish a mapping between a *portType* and *port* in the WSDL file. There can only be one entry for each *portType*.
  - e. Select the new WSDL file.
  - f. Click **OK**.

The `defaultMappings` of a deployed Web service can also be changed using the administrative console. Using application management, navigate to the Web module or EJB module of the Web service application and select **Web Services Client Bindings**.

9. Access the **Web Services Client Port Bindings** editor through the **Port Bindings** tab at the bottom of the editor pane.
10. Verify the `syncTimeout` element settings.
  - a. Create a **Port Qualified Name Bindings** for the port.
  - b. Open the **Web Services Client Bindings** editor.
  - c. Confirm that a service reference is selected in either the **Component scoped references** or **Service references** section.
  - d. Expand the **Port qualified name bindings** section.
  - e. Click **Add**. A new entry is now added to the **Port qualified name bindings** list.
  - f. Click the new **Port qualified name bindings** entry. The **Web Services Client Port Bindings** editor opens.
  - g. Expand the **Port qualified name bindings details** section.
  - h. Type the *namespace* of the WSDL file port you want to configure, in the **Port Namespace Link** field.
  - i. Type the *local\_name* of the WSDL file port you want to configure in the **Port Local Name Link** field. The name displayed in the **Port qualified name bindings** list is the local name of the WSDL file port.
  - j. Click **OK**.
  - a. Configure the `syncTimeout` property by locating the **Synchronization timeout** field and enter the desired value.
11. Verify the `basicAuth` element settings.
  - a. Open the **Web Services Client Bindings** editor.
  - b. Expand the **Basic authentication** section.
  - c. Type the desired value in the **User ID** and **Password** fields.
  - d. Click **OK**.
12. Verify the `sslConfig` element settings.
  - a. Open the **Web Services Client Bindings** editor.
  - b. Expand the **SSL Configuration** section.
  - c. Type the desired value in the **Name** field.



d. Click **OK**.

13. After editing the properties, type **ctrl-s** on your keyboard to save the changes.

## **ibm-webservicesclient-bnd.xmi assembly properties**

The `ibm-webservicesclient-bnd.xmi` file contains information for the Web services runtime that is WebSphere product-specific.

You can edit these properties using the Assembly Toolkit. See *Configuring the ibm-webservicesclient-bnd.xmi deployment descriptor* for instructions.

### **Assembly properties**

The following user-definable assembly properties are supported:

- **componentNameLink**

Attribute of the `componentScopedRefs` element that specifies the link to the corresponding `<component-scoped-refs>` element in `webservicesclient.xml` file. This property is used only when the Web service client is an EJB.

- **serviceRefLink**

Attribute of the `serviceRefs` element that specifies the link to the `<service-ref-name>` in the `webservicesclient.xml` file.

You can edit this property in the Assembly Toolkit:

- **deployedWSDLFile**

Attribute of the `serviceRefs` element is optional and permits an alternate WSDL file to be used other than that specified in the `<wsdl-file>` element of `webservicesclient.xml` file. If this attribute is specified, the alternate WSDL file must be packaged in the same module and must be compatible with the development WSDL file. The `deployedWSDLFile` property is used to supply a new WSDL file containing a different endpoint URL than the original WSDL file.

- **defaultMappings** element

Identifies which port should be used for a given `portType` when none is explicitly selected by the client. This element has the following attributes: `portTypeNamespace`, `portTypeLocalName`, `portNamespace`, `portLocalName`. These attributes identify which `wsdl:port` should be used for a `wsdl:portType`.

- **syncTimeout**

Attribute of the `portQnameBindings` element that specifies how long, in seconds, to wait for a response from a synchronous call.

- **basicAuth**

Element of the `portQnameBindings` element that can be used to authenticate a service client to the service endpoint, independent of the underlying transport that includes, HTTP, HTTPS, and JMS. Set the user ID and password attributes as needed.

- **sslConfig**

Element of the `portQnameBindings` element that specifies the Secure Sockets Layer (SSL) configuration of an HTTPS outbound request. The name attribute is the name of a SSL configuration entry or alias defined in the SSL Configuration Repertoire.

**Note:** This attribute is only used when the client is running in the WebSphere Application Server.

### **Example bindings file**

The following examples demonstrate the spelling and position of the various attributes. You cannot cut and paste these examples because they do not contain the required ID attributes. If you add elements to a binding file template generated by the **WSDL2Java** command, you must confirm that each element has an ID attribute whose value is a unique string. Review the template `xmi` files generated by the **WSDL2Java** command for examples of ID strings.



```

<com.ibm.etools.webservice.wscbnd:ClientBinding xmi:version="2.0"
xmlns:xmi="http://www.omg.org/XMI" xmlns:com.ibm.etools.webservice.wscbnd=
"http://www.ibm.com/websphere/appserver/schemas/5.0.2/wscbnd.xmi">

 <componentScopedRefs componentNameLink="myComponent ref"/>

 <serviceRefs serviceRefLink="myService ref"
 deployedWSDLFile="META-INF/wsdl/alternate.wsdl">
 <defaultMappings portTypeLocalName="AddressBook"
 portTypeNamespace="http://www.com.ibm"
portLocalName="AddressBookPort" portNamespace="http://www.com.ibm"/>
 <portQnameBindings portQnameNamespaceLink="http://www.com.ibm"
portQnameLocalNameLink="AddressBookPort" syncTimeout="99">
 <basicAuth userid="myId" password="myPassword"/>
 <sslConfig name="mynode/DefaultSSLSettings"/>
 </portQnameBindings>
 </serviceRefs>
</com.ibm.etools.webservice.wscbnd:ClientBinding>

```

## Configuring the webservicescient.xml deployment descriptor

You should have an enterprise JavaBean (EJB) Java archive (JAR) file, Web archive (WAR) file or an application client file that you can import into the Assembly Toolkit.

This topic explains how to configure the webservicescient.xml deployment descriptor with the Assembly Toolkit. It is one of the tools available with the Application Server Toolkit product. For more information about completing tasks with the Assembly Toolkit, click **Help > Help** in the Assembly Toolkit graphical user interface (GUI).

To configure the webservicescient.xml deployment descriptor with the Assembly Toolkit:

1. Start the Assembly Toolkit.
2. Click **File > Import** to import the EJB JAR file, WAR file or application client file into the Assembly Toolkit.
3. Open the J2EE perspective by clicking **Windows > Open Perspective > J2EE**.
4. Switch to the **Project Navigator** pane by clicking the **Project Navigator** tab.
5. Locate the project containing the webservicescient.xml file in the **Project Navigator** pane.
6. Expand the directories under the project until the META-INF or WEB-INF directory and its contents appear.
7. Right-click on the webservicescient.xml file.
8. Select **Open. The Web Services Client** editor opens.
9. Expand the **Service references** section.
10. Select the *service\_reference* that you want to configure.
11. Expand the **Service reference overview** section.
12. Type the name of the service for which the client accesses in the **Description** field.
13. Expand the **Service reference implementation details** section.
  - a. Type the name that the Java Naming Directory Interface (JNDI) uses to locate the service in the **Service references name** field. The JNDI lookup string for this service is `java:comp/env/service-ref-name`. By convention, the service reference name always begins with `service/`.
  - b. Type the class name, including package, of the generated Java interface that is the Service Interface for this Web service in the **Service interface name** field.
  - c. Type the WSDL file name used by the client, relative to the root of the module, in the **WSDL file** field.

- d. Type the file name of the Java mapping file, relative to the root of the module, in the **JAX RPC mapping file** field.
14. Click **ctrl-s** to save the changes.

The `webservicesclient.xml` deployment descriptor is configured.

## Configuring the `webservicesclient.xml` deployment descriptor for Handler classes

This topic explains how to use the Assembly Toolkit to configure the `webservicesclient.xml` deployment descriptor for user-provided Handler classes. The Assembly Toolkit is a component of the Application Server Toolkit. For more information about completing tasks with the Assembly Toolkit, click **Help > Help** in the Assembly Toolkit graphical user interface (GUI).

You should have an Enterprise archive (EAR) file for the applications you want to configure. For some handler use, such as logging or tracing, only the server or client application needs to be configured. For other handler use, including sending information in SOAP headers, the client and server applications must be configured with symmetrical handlers.

The modules in the EAR file should contain the handler classes being configured. These classes implement the `javax.xml.rpc.handler.Handler` interface. For more information on writing handler classes, see Chapter 6 of the Web Services for Java 2 platform, Enterprise Edition (J2EE) 1.0 specification and chapter 12 of the Java API for XML-based remote procedure call (JAX-RPC) 1.0 specification available through Web services: Resources for learning. The application modules must contain the `webservices.xml` (for server) and `webservicesclient.xml` (for client) deployment descriptors.

To configure a handler in the `webservicesclient.xml` deployment descriptor:

1. Start the Assembly Toolkit.
2. Click **File > Import** and import the EAR file into the Assembly Toolkit.
3. Open the J2EE perspective by clicking **Windows > Open Perspective > Other > J2EE**.
4. Click the **Project Navigator** tab to switch to the **Project Navigator** pane.
5. Locate the project that contains the `webservicesclient.xml` deployment descriptor. Expand the directories under the project until the `META-INF` or `WEB-INF` directory and its contents, including the `webservicesclient.xml` file, are visible.
6. Right-click the `webservicesclient.xml` file.
7. Click **Open**. The **Service References** pane of the **Web Services Client** editor opens.
8. Expand the **Service references** section.
  - a. Select the *service\_reference* for which you want to configure the handler.
9. Expand the **Handlers** section.
10. Click **Add** at the bottom of the **Handlers** section. A **Class browser** opens.
11. Browse for the name of the Handler class in the module. When it displays in the **Matching types** field, select the class and click **OK**. The Class browser window closes after you click OK and the **Handlers** pane of the **Web Services Editor** opens.
12. (Optional) Configure properties in the **Handlers** pane. See Handler class properties for a list of the properties you can configure in this step.
13. Type **ctrl-s** to save the changes.

### Handler class properties

You can configure the following Handler class properties through the Assembly Toolkit. See Configuring the `webservices.xml` deployment descriptor for Handler classes or Configuring the `webservicesclient.xml` deployment descriptors for Handler classes for instructions on how to configure the properties.

**Description**

Standard Java 2 platform, Enterprise Edition (J2EE) technology descriptor field.

**Display name**

Standard J2EE technology descriptor field.

**Small icon**

Standard J2EE technology descriptor field.

**Large icon**

Standard J2EE technology descriptor field.

**Handler name**

The name of the handler. This name must be unique within the module.

**Handler class**

The fully qualified name of the Handler class. Initially, it is set by the Assembly Toolkit's class browser.

**Initial parameters**

Property names and values to be made available to the handler.

**SOAP headers**

Qnames of the SOAP headers that are processed by this handler. See section 12.2.2 of the Java API for XML-based remote procedure call (JAX-RPC) 1.0 specification, available through Web services: Resources for learning, for more information about setting this property.

**SOAP roles**

URIs containing the SOAP actor names for which the handler acts in the role of. See section 12.2.2 of the Java API for XML-based remote procedure call (JAX-RPC) 1.0 specification, available through Web services: Resources for learning, for more information about setting this property.

**Example: Configuring Handler classes for Web services deployment descriptors**

This scenario explains how to add trivial client and server Handler classes to a sample application named `WebServicesSamples.ear`. The Handler classes display messages when given a request or response to handle.

The code for the client Handler class is:

```
package samples;

public class ClientHandler implements javax.xml.rpc.handler.Handler {
 public ClientHandler() { }

 public boolean handleRequest(javax.xml.rpc.handler.MessageContext
 context) {
 System.out.println("ClientHandler: In handleRequest");
 return true; }

 public boolean handleResponse(javax.xml.rpc.handler.MessageContext
 context) {
```

```

 System.out.println("ClientHandler: In handleResponse");
 return true; }

public boolean handleFault(javax.xml.rpc.handler.MessageContext
context) {
 System.out.println("ClientHandler: In handleFault");
 return true; }

public void init(javax.xml.rpc.handler.HandlerInfo config) { }

public void destroy() {
}

public javax.xml.namespace.QName[] getHeaders() {
 return null; }
}

```

The code for the server Handler class is:

```

package sample;
public class ServerHandler implements javax.xml.rpc.handler.Handler {
 public ServerHandler() { }

public boolean handleRequest(javax.xml.rpc.handler.MessageContext
context) {
 System.out.println("ServerHandler: In handleRequest");
 return true; }

public boolean handleResponse(javax.xml.rpc.handler.MessageContext
context) {
 System.out.println("ServerHandler: In handleResponse");
 return true; }

public boolean handleFault(javax.xml.rpc.handler.MessageContext
context) {
 System.out.println("ServerHandler: In handleFault");
 return true; }

public void init(javax.xml.rpc.handler.HandlerInfo config) { }

public void destroy() { }

public javax.xml.namespace.QName[] getHeaders() {
 return null; }
}

```

1. Compile these classes using
  - %JAVA\_HOME%\bin\java -extdirs %WAS\_EXT\_DIRS% ClientHandler.java ServerHandler.java (on Windows)
  - \$JAVA\_HOME/bin/java -extdirs \$WAS\_EXT\_DIRS ClientHandler.java ServerHandler.java (on Unix)
2. Open the Assembly Toolkit and import the two sample EAR files:
  - %WAS\_HOME%\samples\lib\WebServicesSamples\WebServicesSamples.ear on Windows or \$WAS\_HOME/samples/lib/WebServicesSamples/WebServicesSamples.ear on Unix.
  - %WAS\_HOME%\samples\lib\WebServicesSamples\ApplicationClients.ear on Windows or \$WAS\_HOME/samples/lib/WebServicesSamples/ApplicationClients.ear on Unix..
3. Import the compiled handler classes into the projects for the sample modules:
  - Import sample.ClientHandler into the **appClientModule** directory of the **AddressBookClient** project.
  - Import sample.ServerHandler into the **ejbModule** directory of the **AddressBookW2JE** project.
4. Configure the webservicessclient.xml deployment descriptor for Handler classes.
5. Configure the webservicess.xml deployment descriptor for Handler classes.
6. Save your changes and export the EAR files.

7. Uninstall the WebServicesSamples.ear application from your server if it is already installed.
8. Install the new WebServicesSamples.ear application.
9. Start the server.
10. Run the client:

**launchClient ApplicationClients.ear -CCjar=AddressBookClient.jar**

When the client executes, the console output is as shown below. The messages from the handlers are shown in bold.

```
IBM WebSphere Application Server, Release 5.1
J2EE Application Client Tool
Copyright IBM Corp., 1997-2003
WSCL0012I: Processing command line arguments.
WSCL0013I: Initializing the J2EE Application Client
Environment.
WSCL0035I: Initialization of the J2EE Application Client
Environment has completed.
WSCL0014I: Invoking the Application Client class
com.ibm.websphere.samples.webservices.addr.AddressBookClient
>> Querying address for 'Purdue Boilermaker' using port
AddressBookW2JE
ClientHandler: In handleRequest
ClientHandler: In handleResponse
>> Response is:
 1 University Drive
 West Lafayette, IN 47907
 Phone: (765) 555-4900
>> Querying address for 'Purdue Boilermaker' using port
AddressBookJ2WE
ClientHandler: In handleRequest
ClientHandler: In handleResponse
>> Response is:
 2 University Drive
 West Lafayette, IN 47907
 Phone: (765) 555-4900
>> Querying address for 'Purdue Boilermaker' using port
AddressBookJ2WB
ClientHandler: In handleRequest
ClientHandler: In handleResponse
>> Response is:
 3 University Drive
 West Lafayette, IN 47907
 Phone: (765) 555-4900
>> Querying address for 'Purdue Boilermaker' using port AddressBookW2JB
ClientHandler: In handleRequest
ClientHandler: In handleResponse
>> Response is:
 4 University Drive
 West Lafayette, IN 47907
 Phone: (765) 555-4900
```

For the client, the Handler class is configured for each service reference, not for each port. The AddressBook sample has four ports, but only one service reference, therefore the ClientHandler handles requests and responses on all ports.

When the server log file is examined, it contains:

```
[9/24/03 16:39:22:661 CDT] 4deec1c6 WebGroup I SRVE0180I:
[HTTP router for AddressBookW2JE.jar] [/AddressBookW2JE] [Servlet.LOG]:
AddressBook: init
[9/24/03 16:39:23:161 CDT] 4deec1c6 SystemOut 0 ServerHandler: In handleRequest
[9/24/03 16:39:23:211 CDT] 4deec1c6 SystemOut 0 ServerHandler: In handleResponse
```

## What to do next

Install and test the application.

## Testing Web services-enabled clients

Before testing your Java client, confirm that the server endpoint specified in the client Web Services Description Language (WSDL) file is running and available.

The following steps and examples assume that you are testing a system that has WebSphere Application Server installed, and that you have configured your environment as described in Setting up a Web services development and unmanaged client execution environment.

Tests are run differently depending on whether the client module has client container deployment information, which consists of the `application-client.xml` and `webservicesclient.xml` files, as well as the JAX-RPC mapping file and WSDL file. The client enterprise archive (EAR) files discussed in this topic are referred to as managed because they contain the deployment information. The client Java archive (JAR) files discussed are referred to as unmanaged because they that do not contain the deployment information.

To test Web services-enabled clients:

1. Test an unmanaged client JAR file.

- a. Execute your application with the **java** command. On Windows platforms:

```
%JAVA_HOME%\bin\java" "-Xbootclasspath/p:%WAS_BOOTCLASSPATH%"
-Djava.security.auth.login.config="%WAS_HOME%\properties\wsjaas_client.conf"
-Djava.ext.dirs="%WAS_EXT_DIRS%"
-classpath "%WAS_CLASSPATH%;<list your application JAR files and classes>"
<fully qualified class name to run><your application parameters>
```

On UNIX:

```
`${JAVA_HOME}/bin/java" "-Xbootclasspath/p:$WAS_BOOTCLASSPATH"
-Djava.security.auth.login.config="$WAS_HOME/properties/wsjaas_client.conf"
-Djava.ext.dirs="$WAS_EXT_DIRS"
-classpath "$WAS_CLASSPATH;<list of your application JAR files and classes>"
<fully qualified class name to run><your application parameters>
```

The unmanaged client application runs.

2. Test a managed client EAR file.

- a. Execute your client application with the **launchClient** command. An example of using the command is as follows:

```
launchClient clientEar
```

Web services-enabled clients that have been tested.

Troubleshoot your Web services application.

---

## Assembling Web services applications based on Web Services for J2EE

This topic explains how to assemble a Web services application that is based on the Web Services for Java 2 platform, Enterprise Edition (J2EE) specification.

You can assemble Web Services for J2EE modules with the Assembly Toolkit which replaces the Application Assembly Tool (AAT). The Assembly Toolkit is one of the tools available with the Application Server Toolkit product. .

To assemble Web services applications:

1. Start the Assembly Toolkit.
2. Assemble a Web services-enabled EJB JAR file.
3. Assemble a Web services-enabled EJB JAR file into an EAR file.
4. (Optional) Enable the EAR file. When the EAR file contains EJB modules, it must have the Web services endpoint WAR file added with the **endptEnabler** command-line tool or Assembly Toolkit before deployment.
5. Assemble a Web services-enabled WAR file.
6. Assemble a Web services-enabled WAR file into an EAR file.

A Web services-enabled EAR file that you can deploy into WebSphere Application Server.

Deploy the Web services-enabled EAR file into WebSphere Application Server.

## Assembling a Web services-enabled EJB JAR file

You can assemble a Web services-enabled enterprise JavaBean (EJB) Java archive (JAR) file in one of two ways:

1. Assemble a Web services-enabled EJB JAR file when starting from Java code.
2. Assemble a Web services-enabled EJB JAR file when starting from Web Services Description Language (WSDL).

An assembled Web services-enabled EJB JAR file.

Configure the webservices.xml deployment descriptor .

## Assembling a Web services-enabled EJB JAR file when starting from Java code

You need the following artifacts:

- Assembled Enterprise JavaBean (EJB) Java archive (JAR) file (not enabled for Web services)
- Compiled Java class for the Service Endpoint Interface
- Web Services Description Language (WSDL) file
- Complete `webservices.xml`, `ibm-webservices-bnd.xmi`, `ibm-webservices-ext.xmi` and Java API for XML-based remote procedure call (JAX-RPC) mapping deployment descriptors.

This topic explains how to assemble a Web service-enabled EJB JAR file with the Assembly Toolkit. The Assembly Toolkit replaces the Application Assembly Tool (AAT) and is one of the tools available with the Application Server Toolkit product.

To assemble an Web services-enabled EJB JAR file when starting from Java code:

1. Start the Assembly Toolkit.
2. Click **File > Import** to import the EJB JAR file into the Assembly Toolkit.
3. Open the J2EE perspective by clicking **Windows > Open Perspective > Other > J2EE**.
4. Switch to the **Project Navigator** pane by clicking the **Project Navigator** tab.
5. Locate the project containing the JAR file you just imported in the **Project Navigator** pane.
6. Expand the `ejbModule` entry until the `META-INF` directory displays. Expand the `META-INF` directory.
7. Right-click the `META-INF` directory and click **New > Folder**. Create a subfolder named `wsdl` in the `META-INF` directory.
8. Copy the WSDL file to the `META-INF\wsdl` directory by right-clicking on the `wsdl` directory and click **File > Import > File system**. Browse the WSDL file for this Web service and click **Finish**.
9. Copy the JAX-RPC mapping file, `webservices.xml`, `ibm-webservices-bnd.xmi`, and `ibm-webservices-ext.xmi` files into the `META-INF` directory.



10. Import the Service Endpoint Interface class so its package begins in the `ejbModule` directory. You can import either the source file or compiled class file. If you import the source file it automatically compiles.

The artifacts required to Web service-enable an EJB module for Web services are added to the JAR file.

After assembling a JAR file named `AddressBook.jar`, the JAR file contains the following files. The files added in this task are in bold:

```
META-INF/MANIFEST.MF
META-INF/ejb-jar.xml
addr/Address.class
addr/AddressBook_RI.class
addr/AddressBookBean.class
addr/AddressBookHome.class
addr/Phone.class
addr/StateType.class
addr/AddressBook.class
META-INF/wsd1/AddressBook.wsd1
META-INF/ibm-webservices-bnd.xmi
META-INF/ibm-webservices-ext.xmi
META-INF/webservices.xml
META-INF/AddressBook_mapping.xml
```

Configure the `webservices.xml` deployment descriptor .

## Assembling Web services-enabled EJB JAR file when starting from WSDL

You need the following artifacts:

- An assembled Enterprise JavaBean (EJB) Java archive (JAR) file containing the EJB implementation and all classes generated by the **WSDL2Java** command tool when the **role** argument is `develop-server` and the **container** argument is `EJB`.
- A Web Services Description Language (WSDL) file
- Complete `webservices.xml`, `ibm-webservices-bnd.xmi`, `ibm-webservices-ext.xmi`, and Java API for XML-based remote procedure call (JAX-RPC) mapping deployment descriptors.

You can use the Assembly Toolkit to assemble Web services-enabled JAR files. The Assembly Toolkit replaces the Application Assembly Tool (AAT). It is one of the tools available with the Application Server Toolkit product.

To assemble a Web services-enabled EJB JAR file when starting from WSDL:

1. Start the Assembly Toolkit.
2. Click **File > Import** to import the EJB JAR file into the Assembly Toolkit.
3. Open the J2EE perspective by clicking **Windows > Open Perspective > Other > J2EE**.
4. Switch to the **Project Navigator** pane by clicking the **Project Navigator** tab.
5. Locate the project for the JAR file you just imported in the **Project Navigator** pane.
6. Expand the `ejbModule` entry so the `META-INF` directory is displayed. Expand the `META-INF` directory.
7. Right-click the `META-INF` directory and select **New > Folder**. Create a subfolder named `wsdl` in the `META-INF` directory.
8. Copy the WSDL file to the `META-INF\wsdl` directory by right-clicking on the `wsdl` directory and click **File > Import > File system**. Browse the WSDL file for this Web service and click **Finish**.
9. Copy the JAX-RPC mapping file as specified by the deployment descriptor `<jaxrpc-mapping-file>` element of `webservices.xml`.
10. Copy `webservices.xml`, `ibm-webservices-bnd.xmi` and `ibm-webservices-ext.xmi` into the `META-INF` subdirectory in the same manner.

The artifacts required to enable an EJB module for Web services are added to the JAR file.

After assembling a JAR file named `AddressBook.jar` contains the following files. The files added in this task are in bold:

```
META-INF/MANIFEST.MF
META-INF/ejb-jar.xml
addr/Address.class
addr/AddressBook_RI.class
addr/AddressBookSoapBindingImpl.class
addr/AddressBookHome.class
addr/Phone.class
addr/StateType.class
addr/AddressBook.class
META-INF/wsdl/AddressBook.wsdl
META-INF/ibm-webservices-bnd.xmi
META-INF/ibm-webservices-ext.xmi
META-INF/webservices.xml
META-INF/AddressBook_mapping.xml
```

Configure the `webservices.xml` deployment descriptor .

## Assembling a Web services-enabled WAR file

You can assemble a Web services-enabled Web archive (WAR) file in one of two ways:

1. Assemble a Web services-enabled WAR file when starting from Java code.
2. Assemble a Web services-enabled WAR file when starting from WSDL.

A Web services-enabled WAR file is assembled.

### Assembling a Web services-enabled WAR file when starting from Java code

You need the following artifacts:

- An assembled Web archive (WAR) file containing `web.xml`, but not Web services-enabled
- The Java class for the Service Endpoint Interface
- A Web Services Description Language (WSDL) file
- Complete `webservices.xml`, `ibm-webservices-bnd.xmi`, `ibm-webservices-ext.xmi`, and Java API for XML-based remote procedure call (JAX-RPC) mapping deployment descriptors.

You can use the Assembly Toolkit to assemble a Web services-enabled WAR file. The Assembly Toolkit replaces the Application Assembly Tool (AAT). It is one of the tools available with the Application Server Toolkit product.

To assemble a Web services-enabled WAR file when starting from Java code:

1. Start the Assembly Toolkit.
2. Click **File > Import** to import the WAR file into the Assembly Toolkit.
3. Open the J2EE perspective by clicking **Windows > Open Perspective > Other > J2EE**.
4. Switch to the **Project Navigator** pane by clicking the **Project Navigator** tab.
5. Locate the project for the WAR file you just imported in the **Project Navigator** pane.
6. Expand the `WebContent` directory so the `WEB-INF` directory is displayed. Expand the `WEB-INF` directory
7. Confirm that the `WEB-INF/web.xml` descriptor for the Web module contains a `<servlet-class>` element indicating the Java bean class that is implementing the service. Confirm by:
  - a. Double-click **Web Deployment Descriptor**.
  - b. In the **Web Deployment Descriptor** editor, click the **Servlets**.
  - c. Enter the full path name of the Java bean class implementing the Web service in the **Servlet class** field.
  - d. Close the editor window to save your changes.

8. Right-click the WEB-INF directory and click **New > Folder**. Create a subfolder named wsd1 in the WEB-INF directory.
9. Copy the WSDL file to the WEB-INF\wsd1 directory by right-clicking on the wsd1 directory and click **File > Import > File system**. Browse the WSDL file for this Web service and click **Finish**.
10. Copy the JAX-RPC mapping file as specified by the deployment descriptor <jaxrpc-mapping-file> element of webservices.xml.
11. Copy webservices.xml,ibm-webservices-bnd.xmi and ibm-webservices-ext.xmi into the WEB-INF subdirectory in the same manner.
12. Import the Service Endpoint Interface class so that its package begins in the JavaSource directory. When you import the source file it is automatically compiled.

The artifacts required to Web service-enable the Web module are added to the WAR file.

Assemble a Web services-enabled WAR into an EAR file.

## Assembling a Web services-enabled WAR file when starting from WSDL

You need the following artifacts:

- Assembled Web archive (WAR) file containing the enterprise JavaBean (EJB) implementation, all classes generated by the **WSDL2Java** command tool and a Web deployment descriptor, web.xml.
- A Web Services Description Language (WSDL) file
- Complete webservices.xml, ibm-webservices-bnd.xmi, ibm-webservices-ext.xmi, and Java API for XML-based remote procedure call (JAX-RPC) mapping deployment descriptors.

You can use the Assembly Toolkit to assemble Web services-enabled WAR files. The Assembly Toolkit replaces the Application Assembly Tool (AAT). It is one of the tools available with the Application Server Toolkit product.

To assemble a Web services-enabled WAR file when starting from WSDL:

1. Start the Assembly Toolkit.
2. Click **File > Import** to import the WAR file into the Assembly Toolkit.
3. Open the J2EE perspective by clicking **Windows > Open Perspective > Other > J2EE**.
4. Switch to the **Project Navigator** pane by clicking the **Project Navigator** tab.
5. Locate the project for the WAR file you just imported in the **Project Navigator** pane.
6. Expand the WebContent directory so the WEB-INF directory is displayed. Expand the WEB-INF directory
7. Confirm that the WEB-INF/web.xml deployment descriptor for the Web module contains a <servlet> element including the <servlet-name> element. To confirm:
  - a. Double-click **Web Deployment Descriptor**.
  - b. In the **Web Deployment Descriptor** editor click the **Servlets** tab.
  - c. Enter the full path name of the Java bean class implementing the Web service in the **Servlet class** field.
  - d. Close the editor window to save your change.
8. Right-click the WEB-INF directory and select **New > Folder**. Create a subfolder named wsd1 in the WEB-INF directory.
9. Copy the WSDL file to the WEB-INF\wsd1 directory by right-clicking on the wsd1 directory and click **File > Import > File system**. Browse the WSDL file for this Web service and click **Finish**.
10. Copy the JAX-RPC mapping file as specified by the deployment descriptor <jaxrpc-mapping-file> element of webservices.xml.
11. Copy the webservices.xml, ibm-webservices-ext.xmi, ibm-webservices-bnd.xmi deployment descriptors in the WEB-INF subdirectory.

The artifacts required to Web service-enable the Web module is added to the WAR file.

Assemble a Web services-enabled WAR into an EAR file.

## Assembling a Web services-enabled EJB JAR into an EAR file

Before assembling a Web services-enabled enterprise archive (EAR) file Assemble a Web services-enabled EJB Java archive (JAR) file.

You can assemble a Web services-enabled EAR file with the Assembly Toolkit. The Assembly Toolkit replaces the Application Assembly Tool (AAT). It is one of the tools available with the Application Server Toolkit product.

To assemble a Web services-enabled EAR file:

1. Start the Assembly Toolkit.
2. Assemble the Web services-enabled JAR file into an EAR file. The EAR file can contain an enterprise bean or application client JAR files, WAR files, Web applications, and metadata describing the applications or `application.xml` files.

A Web services-enabled EAR file.

In the following example, there is an `application.xml` deployment descriptor packaged with a Web services-enabled JAR file called `AddressBook.jar` that is packaged into an EAR file called `AddressBook.ear`. The EAR file contains:

```
META-INF/MANIFEST.MF
META-INF/application.xml
AddressBook.jar
```

An example of the `application.xml` deployment descriptor is as follows:

```
<?xml version="1.0" encoding="UTF-8"?>
<!DOCTYPE application PUBLIC "-//Sun Microsystems, Inc.//DTD J2EE Application 1.3//EN"
"http://java.sun.com/dtd/application_1_3.dtd">
<application id="Application_ID">
 <display-name>AddressBookJ2WEE</display-name>
 <description>AddressBook EJB Example from Java</description>
 <module id="EjbModule_1">
 <ejb>AddressBook.jar</ejb>
 </module>
</application>
```

Enable the EAR file. Then, deploy the EAR file into WebSphere Application Server.

## Assembling a Web services-enabled WAR into an EAR file

Before assembling a Web services-enabled enterprise archive (EAR) file Assemble a Web services-enabled Web archive (WAR) file.

This topic explains how to assemble a Web services-enabled WAR file into and EAR file using the Assembly Toolkit. The Assembly Toolkit replaces the Application Assembly Tool (AAT) and is one of the tools available with the Application Server Toolkit product.

To assemble a Web services-enabled WAR file into an EAR file:

1. Start the Assembly Toolkit.

2. Assemble the Web services-enabled WAR file into an EAR file. Now assemble the EAR file that contains the JAR or WAR files. The EAR file can contain an enterprise bean or application client JAR files; Web applications or WAR files; and metadata describing the applications or application.xml files.

A Web services-enabled EAR file.

In the following example, there is an application.xml deployment descriptor packaged with a Web services-enabled JAR file called AddressBook.jar that is packaged into an EAR file called AddressBook.ear. The EAR file contains:

```
META-INF/MANIFEST.MF
META-INF/application.xml
AddressBook.war
```

An example of the application.xml deployment descriptor is as follows:

```
<?xml version="1.0" encoding="UTF-8"?>
<!DOCTYPE application PUBLIC "-//Sun Microsystems, Inc.//DTD J2EE Application 1.3//EN"
"http://java.sun.com/dtd/application_1_3.dtd">
<application id="Application_ID">
 <display-name>AddressBook</display-name>
 <description>AddressBook Example from Java bean</description>
 <module id="WebModule_1">
 <web>
 <web-uri>AddressBook.war</web-uri>
 <context-root>/AddressBook</context-root>
 </web>
 </module>
</application>
```

Deploy Web services based on Web Services for J2EE.

## Enabling a Web services-enabled EAR file

Before doing this task, you need to Assemble a Web services-enabled EJB JAR into an enterprise archive (EAR) file.

You can add router modules to your Web services-enabled application, also known as an EAR file with the **endptEnabler** command or the Assembly Toolkit. The Assembly Toolkit replaces the Application Assembly Toolkit (AAT) and is a component of the Application Server Toolkit (ASTK) product.

These tools add one or more router modules to the EAR file for each EJB JAR module within the EAR file. A router module provides an endpoint for the Web services in a particular enterprise JavaBean (EJB) Java archive (JAR) module.

Each router module supports a specific transport such as HyperText Transport Protocol (HTTP) or Java Messaging Service (JMS). If there are no EJB JAR modules in the EAR file, it is not necessary to use these tools.

1. Enable an EAR file with the **endptEnabler** command-line tool.
2. Enable an EAR file with the Assembly Toolkit.

Deploy the EAR file into WebSphere Application Server.

## Enabling a Web services-enabled EAR file with the endptEnabler command

Before doing this task, you need to assemble a Web services-enabled EJB JAR into an enterprise archive (EAR) file.

The **endptEnabler** command-line tool adds one or more router modules to the EAR file for each EJB JAR module within the EAR file. A router module provides an endpoint for the Web services in a particular Enterprise JavaBean (EJB) Java archive (JAR) module.

Each router module supports a specific transport such as HyperText Transport Protocol (HTTP) or Java Messaging Service (JMS). If there are no EJB JAR modules in the EAR file, it is not necessary to use these tools.

To enable an EAR file with the **endptEnabler** command:

1. Invoke the **endptEnabler** command from the *install\_root*\bin directory. If you are using UNIX, invoke the command from the *install\_root*/bin directory.
2. Enter the name of the EAR file, when prompted.
3. Enter various input values as they are requested by the **endptEnabler** command. You are prompted for various input values for each Web services-enabled EJB JAR module in the EAR file. Typically, you should accept the defaults for each prompt. See **endptEnabler** prompts and commands for more information about **endptEnabler** command prompts.

An HTTP or JMS router module is added to the EAR file for each Web services-enabled EJB JAR module contained in the EAR file. For HTTP, a context-root is configured for the application so the Web service can be invoked through a URL. The URL used to invoke the Web service is:

```
http://host[:port]/context-root/services/port-component-name
```

Deploy the EAR file into WebSphere Application Server.

**endptEnabler command:** The **endptEnabler** command enables a set of Web services within an enterprise archive (EAR) file. You can add one or more router modules to the EAR file that include a Web service-enabled EJB JAR file.

Each router module provides a Web service endpoint for a particular transport. For example, an HyperText Transport Protocol (HTTP) router module can be added so that the Web service can receive requests over the HTTP transport, and a Java Messaging Service (JMS) router module can be added so that the Web service can receive requests from a JMS queue or topic.

In its interactive mode, the **endptEnabler** command guides you through the required steps to enable one or more services within an application. The **endptEnabler** command makes a backup copy of your original EAR file in the event that you need to remove or add services at a later time. If your EAR file contains a Web service-enabled EJB JAR file, you must run the **endptEnabler** command before the EAR file is deployed. Otherwise, you do not need to run the command.

### **endptEnabler usage syntax**

Invoke the **endptEnabler** command from the WebSphere Application Server bin directory. The command syntax is as follows:

```
endptEnabler
 [-verbose|-v]
 [-quiet|-q]
 [-help|-h|-?]
 [-properties|-p properties-filename]
 [-transport|-t default-transport]
 [-enableHttpRouterSecurity]
 [ear-filename]
```

All parameters are optional and described as follows:

- **-verbose, -v**

Detailed progress messages are displayed as the tool processes the EAR file. This command-line option is mapped to the verbose global property.



- **-quiet, -q**

No per-module progress messages are displayed as the tool processes the EAR file. This command-line option is mapped to the quiet global property.

- **-help, -h, -?**

A brief help message is displayed explaining the various options.

- **-properties, -p <properties-filename>**

Properties from the file *properties-filename* are read and used to control the behavior of the tool.

- **-transport, -t <default-transport>**

Specifies the default list of transports for which router modules should be created for each EJB JAR file contained in the EAR file. This command-line option is mapped to the defaultTransports global property.

Examples are:

```
-transport http (the default)
-transport jms
-t http,jms
```

- **-enableHttpRouterSecurity**

Enables you to add a security policy for all authenticated users to protect the HTTP router module if all the EJB's are secured in the EJB JAR file. This command-line option is mapped to the http.enableRouterSecurity global property.

- **<ear-filename>**

Specifies the name of the EAR file to be processed.

If the *ear-filename* parameter is not entered on the command line, the interactive mode is used. In interactive mode, you are prompted for the EAR file name, router module names and other important values as the processing occurs. The following dialog is an example of the endptEnabler interactive mode:

**Note:** In this dialog, user input is in fixed width font, and endptEnabler output is in bold.

```
endptEnabler<enter>
WSWS2004I: IBM WebSphere Application Server Release 5
WSWS2005I: Web Services Enterprise Archive Endpoint Enabler Tool.
WSWS2007I: (C) COPYRIGHT International Business Machines Corp. 1997, 2003
WSWS2006I: Please enter the name of your EAR file: AddressBook.ear<enter>

WSWS2003I: Backing up EAR file to: AddressBook.ear~

WSWS2016I: Loading EAR file: AddressBook.ear
WSWS2017I: Found EJB Module: AddressBookEJB.jar

WSWS2029I: Enter http router name for EJB Module AddressBookEJB
[AddressBookEJB_HTTPRouter.war]:<enter>
WSWS2030I: Enter http context root for EJB Module AddressBookEJB
[/AddressBookEJB]:<enter>
WSWS2024I: Adding http router for EJB Module AddressBookEJB.jar.
WSWS2036I: Saving EAR file AddressBook.ear...
WSWS2037I: Finished saving the EAR file.
WSWS2018I: Finished processing EAR file AddressBook.ear.
```

If the *ear-filename* parameter is entered on the command line, the non-interactive mode is used. In non-interactive mode, router module names and other important values are determined from user-specified properties or default values.

## endptEnabler properties

The **endptEnabler** command allows you to control its run time behavior by specifying a set of properties with the **-properties** command-line option. These properties fall into two categories: global and per-module. Global properties affect the overall behavior of the tool as it processes multiple EJB JAR modules within the EAR file. Per-module properties affect the processing of a particular EJB JAR module.

## Global properties



The following table describes the global properties supported by the **endptEnabler** command:

Property name	Description	Default value
verbose	Displays detailed progress messages.	False
quiet	Displays only brief progress messages.	False
http.enableRouter Security	Enables you to add a security policy for all authenticated users to protect the HTTP router module if all the EJB's are secured in the EJB JAR file.	False
http.router ModuleNameSuffix	Specifies the suffix used to construct default HTTP router module names. The .war extension is added by the <b>endptEnabler</b> command.	_HTTPRouter
jms.routerModule NameSuffix	Specifies the suffix used to construct default JMS router module names. The .jar extension is added by the <b>endptEnabler</b> command.	_JMSRouter
jms.listenerInput PortNameSuffix	Specifies the suffix used to construct default Listener Input Port names.	_ListenerPort
jms.default DestinationType	Specifies the default destination type to use for all JMS router modules added to the EAR file. This should be either queue or topic.	queue
defaultTransports	Specifies the default list of transports for which router modules should be created. The list can contain the values http and jms. Multiple values are separated by a comma. Examples are: http, jms and http,jms.	http

### Per-module properties

The following table describes the per-module properties supported by the **endptEnabler** command. *ejbJarName* refers to the name of an EJB JAR module within the EAR file, without the .jar extension.

Property name	Description	Default value
<ejbJarName> .transports	Lists the transports for which router modules should be created for a particular EJB JAR file. The list can contain the values http and jms. Multiple values are separated by a comma. Examples are: http, jms and http,jms.	http
<ejbJarName>.http.skip	Specifies the flag which bypasses the addition of an HTTP router module even if it would otherwise be added (based on other properties). Valid values are true and false.	False
<ejbJarName> .http.routerModuleName	Specifies the name of the HTTP router module for a particular EJB JAR file.	<i>ejbJarName</i> _HTTPRouter

<ejbJarName>.http.contextRoot	Specifies the context root associated with the HTTP router module for a particular EJB JAR file.	<i>/ejbJarName</i>
<ejbJarName>.jms.skip	Specifies the Flag which bypasses the addition of an HTTP router module even if it would otherwise be added (based on other properties). Valid values are true and false.	false
<ejbJarName>.jms.routerModuleName	Specifies the name of the JMS router module for a particular EJB JAR file.	<i>ejbJarName_JMSRouter</i>
<ejbJarName>.jms.listenerInputPortName	Specifies the name of the Listener Input Port to be associated with the JMS router module.	<i>ejbJarName_ListenerPort</i>
<ejbJarName>.ejb JarName>.jms.destinationType	Specifies the JMS destination type associated with the JMS router. Valid values are queue and topic.	queue

### Properties example

Suppose an EAR file contains an EJB JAR file named, StockQuoteEJB.jar that contains Web services. The following set of properties might be used to control the **endptEnabler** command runtime behavior as it processes the EAR file:

```
StockQuoteEJB.transports=http,jms
```

```
StockQuoteEJB.http.routerModuleName=StockQuoteEJB_HTTP
```

```
StockQuoteEJB.http.contextRoot=/StockQuote
```

```
StockQuoteEJB.jms.routerModuleName=StockQuoteEJB_JMS
```

```
StockQuoteEJB.jms.listenerInputPortName=StockQuote_LP
```

```
StockQuoteEJB.jms.destinationType=queue
```

### endptEnabler examples

The following commands are examples of how the **endptEnabler** command can be used:

```
endptEnabler MyApp.ear
```

```
endptEnabler -t jms,http MyApp.ear
```

```
endptEnabler -v -properties MyApp.props MyApp.ear
```

```
endptEnabler -q -t jms MyApp.ear
```

### Enabling a Web services-enabled EAR file with the Assembly Toolkit

Before doing this task, you need to Assemble a Web services-enabled EJB JAR into an enterprise archive (EAR) file.

You can add one or more router modules to your Web services-enabled application, also known as an EAR file with the Assembly Toolkit. The Assembly Toolkit replaces the Application Assembly Toolkit (AAT) and is a component of the Application Server Toolkit (ASTK) product.

A router module provides an endpoint for the Web services in a particular Enterprise JavaBean (EJB) Java archive (JAR) module.

Each router module supports a specific transport such as HyperText Transport Protocol (HTTP) or Java Messaging Service (JMS). If there are no EJB JAR modules in the EAR file, it is not necessary to use these tools.

To enable a Web services-enabled EAR file with the Assembly Toolkit:

1. Start the Assembly Toolkit.
2. Right-click on the EJB project to be enabled.
3. Click **Web Services > Endpoint Enabler**.
4. Specify the transport and router module names in the corresponding fields.
5. Click **OK**.

An HTTP or JMS router module is added to the EAR file for each Web services-enabled EJB JAR module contained in the EAR file. For HTTP, a context-root is configured for the application so the Web service can be invoked through a URL. The URL used to invoke the Web service is:

```
http://host[:port]/context-root/services/port-component-name
```

Deploy the EAR file into WebSphere Application Server.

---

## Configuring the webservices.xml deployment descriptor

Create an enterprise JavaBean (EJB) Java archive (JAR) file or Web archive (WAR) file containing `webservices.xml`:

- Assemble a Web services-enabled EJB JAR file when starting from Java code.
- Assemble a Web services-enabled EJB JAR file from WSDL.
- Assemble a Web services-enabled WAR file when starting from Java code.
- Assemble a Web services-enabled WAR file when starting from WSDL.

Do one of the following based on whether your implementation is an EJB JAR file or Web module WAR file:

- Develop Web services Java bean deployment descriptor templates from a WSDL file.
- Develop Web services EJB deployment descriptor templates from a WSDL file.

This topic explains how to configure the `webservices.xml` deployment descriptor with the Assembly Toolkit. It is one of the tools available with the Application Server Toolkit product. For more information about completing tasks with the Assembly Toolkit, click **Help > Help** in the Assembly Toolkit graphical user interface (GUI).

To configure the `webservices.xml` deployment descriptor:

1. Start the Assembly Toolkit.
2. Click **File > Import** to import the EJB JAR file or WAR file into the Assembly Toolkit.
3. Open the J2EE perspective by clicking **Windows > Open Perspective > J2EE**.
4. Switch to the **Project Navigator** pane by clicking the **Project Navigator** tab.
5. Locate the project containing the `webservicesclient.xml` file in the **Project Navigator** pane.
6. Expand the directories under the project until the META-INF or WEB-INF directory and its contents appear.
7. Right-click the `webservices.xml` file.
8. Select **Open**. The **Web Services** editor opens.
9. Expand the **Web service descriptions** section.
  - a. Select the service you want to configure.
10. Expand the **Web service description implementation details** section.

- a. Verify the **Web service description name** field is unique among all the Web service descriptions in the editor.
  - b. Verify that the **WSDL file** field indicates there is an existing WSDL file in the module. This file, by convention, should be located in the META-INF/wsdl directory for an enterprise bean JAR file and in the WEB-INF/wsdl directory for a WAR file.
  - c. Verify the **JAX-RPC mapping file** field indicates an existing mapping file within the module. This file, by convention, should be located in the META-INF directory for an enterprise bean JAR file and in the WEB-INF directory for a WAR file.
11. Expand the **Port components** section.
    - a. Verify there are port component entries corresponding to the used WSDL ports in the **Port components** section.
  12. Select a *port\_component* to open the editor for that port component. The **Port Components** editor opens.
  13. Expand the **Port component implementation details** section.
    - a. Verify the **WSDL Port Namespace URL** and **WSDL Port Local part** fields are set to the namespace and local name of the corresponding port in the WSDL file. These fields are configured by the **WSDL2Java** command tool when the webservices.xml file is generated.
  14. Verify the **Service endpoint interface** field names the fully qualified Service Endpoint Interface class. This field is configured by the **WSDL2Java** command when the webservices.xml file is generated.
  15. Locate the **Service implementation bean** field.
    - a. Configure this field to indicate the EJB or servlet that implements the Web service. Configure by selecting **EJB link** for an enterprise bean module or **Servlet link** for a Web module. Use the drop down list in the **Service implementation bean** field to select the enterprise bean or servlet used to implement the Web service. The choices in the drop down menu come from the enterprise beans defined in the ejb-jar.xml file for an enterprise bean module or the servlets defined in the web.xml file for a Web module.

---

## Configuring the ibm-webservices-bnd.xmi deployment descriptor

Develop implementation templates and bindings for the `ibm-webservices-bnd.xmi` from the Web Services Description Language (WSDL) file.

Do one of the following based on whether your implementation is an EJB Java archive (JAR) file or Web module Web archive (WAR) file:

- Assemble a Web services-enabled JAR file when starting from Java code.
- Assemble a Web services-enabled WAR file when starting from Java code.
- Assemble a Web services-enabled JAR file when starting from WSDL.
- Assemble a Web services-enabled WAR file when starting from WSDL.

This topic explains how to configure bindings using the Assembly Toolkit which replaces the Application Assembly Tool (AAT). It is one of the tools available with the Application Server Toolkit product. For more information about completing tasks with the Assembly Toolkit, click **Help > Help** in the Assembly Toolkit graphical user interface. .

To configure the `ibm-webservices-bnd.xmi` deployment descriptor with the Assembly Toolkit:

1. Start the Assembly Toolkit.
2. Click **File > Import** to import the EJB JAR file or WAR file into the Assembly Toolkit.
3. Open the J2EE perspective by clicking **Windows > Open Perspective > Other > J2EE**.
4. Switch to the **Project Navigator** pane by clicking the **Project Navigator** tab.
5. Locate the project containing the `webservices.xml` file in the **Project Navigator** pane.

6. Expand the directories under the project until the META-INF or WEB-INF directory and its contents appear.
7. Right-click the `webservices.xml` file.
8. Select **Open**. The **Web Services** editor opens.
9. Click the **Bindings** tab located at the bottom of the editor pane. The **Web Services Bindings** editor opens.
10. Verify the `wsdescNameLink` element settings.
  - a. Expand the **Web services description bindings** section. Verify that the value of the `<webservice-description-name>` element in the `webservices.xml` deployment descriptor is listed in the section. If the value is not listed:
  - b. (Optional) Click **Add**, choose the correct Web services name and click **OK**. You do not need to complete this step if you have verified that the correct Web services name is listed in the **Web Services Description Bindings** tab.
11. Verify the `pcnameLink` attribute settings.
  - a. Expand the **Web Service Description Bindings** section. Verify that the correct service is selected. If the correct service is not listed:
  - b. (Optional) Expand **Port Component Binding**. Verify the correct Web services name is selected in the **Web Service Description Bindings** section.  
This selection is a prerequisite to creating a `pcnameLink` attribute.
  - c. In the **Port Component Binding** section, click **Add**. You need to make a selection in the **Web Service Description Bindings** section before you can create the port component binding in the **Port Component Binding** section. The **Port Component Bindings Dialog** opens.
  - d. Select the desired port from the drop down list in the **PC Name Link** field.
  - e. Click **OK**.
  - f. Click the **Binding Configurations** tab to view the bindings for your port.
  - g. **Save** the bindings file.
12. Click **File > Export** to export the JAR file, or continue using the Assembly Toolkit for configuration and assembly tasks.
13. Click **ctrl-s** to save your changes.

The `ibm-webservices-bnd.xmi` deployment descriptor is configured for the Web service implementation module.

## ibm-webservices-bnd.xmi assembly properties

### ibm-webservices-bnd.xmi properties

The `ibm-webservices-bnd.xmi` file is a deployment descriptor for a Web Services-enabled Web module or enterprise JavaBean (EJB) module. It contains information for the Web services runtime that is either WebSphere product-specific or was not specified by the Web Services for Java 2 platform, Enterprise Edition (J2EE) specification.

You can edit these properties using the Assembly Toolkit. See *Configuring the `ibm-webservices-bnd.xmi` deployment descriptor* for instructions.

The following user-definable assembly properties are supported:

- **wsDescNameLink**  
Attribute of the `wsdescBindings` element that specifies the link to the corresponding `<webservice-description-name>` in `webservices.xml`.
- **pc-name-link**  
Attribute of the `pcBindings` element that specifies the link to the `<port-component-name>` in the `webservices.xml` file.

- **scope**

Attribute of the `pcBindings` element that specifies when new instances of implementation beans are created. Possible values are Request, Session, and Application.

The value of `scope` for a deployed Web service can be changed using the administrative console. Using application management, navigate to the Web module of the Web service application and select Web Services Implementation Scope.

### Example bindings file

The following examples demonstrate the spelling and position of the various attributes. You cannot cut and paste these examples because they do not contain the required ID attributes. If you add elements to a binding file template generated by the **WSDL2Java** command, you must confirm that each element has an ID attribute whose value is a unique string. Review the template xmi files generated by the **WSDL2Java** command for examples of ID strings.

```
<com.ibm.etools.webservice.wsbind:WSBinding xmi:version="2.0" xmlns:xmi="http://www.omg.org/XMI" xmlns:com.ibm.etools.webservice.wsbind="http://www.ibm.com/websphere/appserver/schemas/5.0.2/wsbind.xmi">
 <wsdescBindings wsDescNameLink="AddressBookService">
 <pcBindings pcNameLink="AddressBook" scope="Application"/>
 </wsdescBindings>
</com.ibm.etools.webservice.wsbind:WSBinding>
```

---

## Configuring the `webservices.xml` deployment descriptor for Handler classes

This topic explains how to use the Assembly Toolkit to configure the `webservices.xml` deployment descriptor for user-provided Handler classes. The Assembly Toolkit is a component of the Application Server Toolkit. For more information about completing tasks with the Assembly Toolkit, click **Help > Help** in the Assembly Toolkit graphical user interface (GUI).

You should have an enterprise archive (EAR) file for the applications you want to configure. For some handler use, such as logging or tracing, only the server or client application needs to be configured. For other handler use, including sending information in SOAP headers, the client and server applications must be configured with symmetrical handlers.

The modules in the EAR file should contain the handler classes being configured. These classes implement the `javax.xml.rpc.handler.Handler` interface. For more information on writing handler classes, see Chapter 6 of the Web Services for J2EE 1.0 specification and chapter 12 of the JAX-RPC 1.0 specification available through Web services: Resources for learning. The application modules must contain the `webservices.xml` (for server) and `webservicesclient.xml` (for client) deployment descriptors.

To configure a handler in the `webservices.xml` deployment descriptor:

1. Start the Assembly Toolkit.
2. Click **File > Import** and import the EAR file into the Assembly Toolkit.
3. Open the J2EE perspective by clicking **Windows > Open Perspective > Other > J2EE**.
4. Click the **Project Navigator** tab to switch to the **Project Navigator** pane.
5. Locate the project that contains the `webservices.xml` deployment descriptor. Expand the directories under the project until the META-INF or WEB-INF directory and its contents, including the `webservices.xml` file, are visible.
6. Right-click the `webservices.xml` file.
7. Click **Open**. The **Web Services** editor opens.
8. Expand the **Web services descriptions** section.



- a. Select the *service* for which you want to configure the handler.
9. Expand the **Port components** section.
10. Select a *port\_component* for which you want the editor to open. The **Port Components** editor opens.
11. Expand the **Port component handlers** section.
12. Click **Add** at the bottom of the **Port component handlers** section. A **Class browser** opens.
13. Browse for the name of the Handler class in the module. When it displays in the **Matching types** field, select the class and click **OK**. The Class browser window closes after you click OK and the **Handlers** pane of the **Web Services Editor** opens.
14. (Optional) Configure properties in the **Handlers** pane. See Handler class properties for a list of the properties you can configure in this step.
15. Type **ctrl-s** to save the changes.

---

## Deploying Web services based on Web Services for J2EE

To deploy Web services that are based on the Web Services for Java 2 platform, Enterprise Edition (J2EE) specification, you need an enterprise application, also known as an enterprise archive (EAR) file that has been configured and enabled for Web services. You can use either the administrative console or the **wsadmin** scripting interface to deploy an EAR file.

If you are installing an application containing Web services by using the **wsadmin** command, specify the **-deployws** option. If you are installing an application containing Web services by using the administrative console, select **Deploy WebServices** during step 1 of the Install New Application wizard. For more information about installing applications using the administrative console see Installing a new application.

### Note:

If the Web services in the application is previously deployed with the **wsdeploy** command, it is not necessary to specify Web services deployment during installation.

Use the following steps to deploy the EAR file with the **wsadmin** command:

1. Start `install_root\bin\wsadmin` from a command prompt. If you are using UNIX start `install_root/bin/wsadmin`.
2. Enter the **\$AdminApp install EARfile "-usedefaultbindings -deployws"** command at the **wsadmin** prompt.

The Web service is installed into the application server.

Secure Web services.

## wsdeploy command

This topic explains how to use the **wsdeploy** command-line tool with Web services that are based on the Web Services for J2EE specification. The **wsdeploy** command adds WebSphere product-specific deployment classes to a Web services compatible enterprise application enterprise archive (EAR) file or an application client Java archive (JAR) file. These classes include:

- Stubs
- Serializers and deserializers
- Implementations of service interfaces

This deployment step must be performed at least once, and can be performed more than once. Deployment can be performed separately using the **wsdeploy** command, the Assembly Toolkit, or when the application is installed. The Assembly Toolkit replaces the Application Assembly Tool (AAT). It is one of the tools available with the Application Server Toolkit product. When using the **wsadmin** command for installation, specify the **-deployws** option. When using the administrative console for installation, select the



**Deploy Web services** check box. When using the Assembly Toolkit, **Right-click** the module and select **Web Services >Deploy Web Services** from the pop-up menu. You can download the Assembly Toolkit from the Web site

[http://www-1.ibm.com/support/docview.wss?rs=180&context=SSEQTP&q=ASTK&uid=swg24005125&loc=en\\_US&cs=utf-8&lang=en+en](http://www-1.ibm.com/support/docview.wss?rs=180&context=SSEQTP&q=ASTK&uid=swg24005125&loc=en_US&cs=utf-8&lang=en+en)

The **wsdeploy** command operates as follows:

- Each module in the enterprise application or JAR file is examined
- If the module contains Web services implementations, indicated by the presence of the `webservicexml` deployment descriptor, the associated Web Services Description Language (WSDL) files are located and the **WSDL2Java** command is run with the role `deploy-server`.
- If the module contains Web services clients, indicated by the presence of the `webservicexml` deployment descriptor, the associated WSDL files are located and the **WSDL2Java** command is run with the role `deploy-client`.
- The files generated by the **WSDL2Java** command are compiled and repackaged.

See `WSDL2Java` command for more information about the files that are generated for deployment.

When the generated files are compiled, they can reference application-specific classes outside the EAR or JAR file if the EAR or JAR file is not self-contained. In this case, use either the `-jardir` or `-cp` option to specify additional JAR or zip files to be added to `CLASSPATH` when the generated files are compiled.

### **wsdeploy command syntax**

The command syntax is as follows:

```
wsdeploy Input_filename Output_filename [options]
```

#### **Required options:**

- ***Input\_filename***  
Specifies the path to the EAR or JAR file to be deployed.
- ***Output\_filename***  
Specifies the path of the deployed EAR or JAR file. If *output\_filename* already exists, it is silently overwritten. The *output\_filename* can be the same as the *input\_filename*.

#### **Other options:**

- **`-jardir` *directory***  
Specifies a directory containing JAR or zip files. All JAR and zip files in this directory are added to the `CLASSPATH` used to compile the generated files. This option can be specified zero or more times.
- **`-cp` *entries***  
Specifies entries to be added to `CLASSPATH` when the generated classes are compiled. Multiple entries are separated the same as they would be in the `CLASSPATH` environment variable, with a semicolon on Windows platforms and a colon for UNIX platforms.
- **`-codegen`**  
Specifies that deployment code is to be generated, but not compiled. This option implicitly specifies the `-keep` option.
- **`-debug`**  
Includes debugging information when compiling, that is, use `javac -g` to compile.
- **`-help`**  
Displays a help message and exit.
- **`-ignoreerrors`**  
Do not stop deployment if validation or compilation errors are encountered.
- **`-keep`**  
Do not delete working directories containing generated classes. A message is displayed indicating the name of the working directory that is retained.

- **-novalidate**  
Do not validate the Web services deployment descriptors in the input file.
- **-trace**  
Displays processing information, including the names of the generated files.

### Example

```
wsdeploy x.ear x_deployed.ear -trace -keep
Processing web service module x_client.jar.
Keeping directory: f:\temp\Base53383.tmp for module: x_client.jar.
Parsing XML file:f:\temp\Base53383.tmp\WarDeploy.wsdl
Generating f:\temp\Base53383.tmp\generatedSource\com\test\WarDeploy.java
Generating f:\temp\Base53383.tmp\generatedSource\com\test\WarDeployLocator.java
Generating f:\temp\Base53383.tmp\generatedSource\com\test\HelloWsBindingStub.java
Compiling f:\temp\Base53383.tmp\generatedSource\com\test\WarDeploy.java.
Compiling f:\temp\Base53383.tmp\generatedSource\com\test\WarDeployLocator.java.
Compiling f:\temp\Base53383.tmp\generatedSource\com\test\HelloWsBindingStub.java.
Done processing module x_client.jar.
```

### Messages

- Flag *-f* is not valid  
Option *f* was not recognized as being a valid option.
- Flag *-c* is ambiguous  
Options may be abbreviated, but the abbreviation must be unique. In this case, the **wsdeploy** command can not determine which option was intended.
- Flag *-c* is missing parameter *-p*  
A required parameter for an option was omitted.
- Missing *p* parameter  
A required option was omitted.

---

## Securing Web services based on WS-Security

Web services security for WebSphere Application Server is based on standards included in the Web services security (WS-Security) specification. These standards address how to provide protection for messages exchanged in a Web service environment. The specification defines the core facilities for protecting the integrity and confidentiality of a message and provides mechanisms for associating security-related claims with the message. Web services security is a message-level standard based on securing Simple Object Access Protocol (SOAP) messages through XML digital signature, confidentiality through XML encryption, and credential propagation through security tokens.

Use the deprecated "Securing Apache SOAP Web services" topics in the WebSphere Application Server, Version 5 documentation if you are still using Apache SOAP Version 2.3.

To secure Web services, you must consider a broad set of security requirements, including authentication, authorization, privacy, trust, integrity, confidentiality, secure communications channels, federation, delegation, and auditing across a spectrum of application and business topologies. One of the key requirements for the security model in today's business environment is the ability to interoperate between formerly incompatible security technologies (such as public key infrastructure, Kerberos and so on.) in heterogeneous environments (such as Microsoft .NET and Java 2 Platform, Enterprise Edition (J2EE)). The complete Web services security protocol stack and technology roadmap is described in Security in a Web Services World: A Proposed Architecture and Roadmap.

Specification: Web Services Security (WS-Security) proposes a standard set of SOAP extensions that you can use to build secure Web services. These standards confirm integrity and confidentiality, which are generally provided with digital signature and encryption technologies. In addition, Web services security provides a general purpose mechanism for associating security tokens with messages. A typical example

of the security token is a user name and password token, in which a user name and password are included as text. Web services security defines how to encode binary security tokens using methods such as X.509 certificates and Kerberos tickets.

To establish a managed environment and to enforce constraints for Web services security, you must perform a Java Naming and Directory Interface (JNDI) lookup on the client to resolve the service reference. For more information on the recommended client programming model, see "Service lookup" in the Java Specification Request (JSR) 109 specification available at: [ftp://www-126.ibm.com/pub/jsr109/spec/1.0/websvcs-1\\_0-fr.pdf](ftp://www-126.ibm.com/pub/jsr109/spec/1.0/websvcs-1_0-fr.pdf).

An administrator can use any of the following methods to integrate message-level security into a WebSphere Application Server environment:

- Secure Web services using XML digital signature
- Secure Web services using XML encryption
- Secure Web services using basicauth authentication
- Secure Web services using identity assertion authentication
- Secure Web services using signature authentication
- Secure Web services using a pluggable token

## Web services security specification- a chronology

This document describes the process used to develop the Web services security specifications.

### Non-OASIS activities

In April 2002, IBM, Microsoft, and VeriSign proposed the *Web Services Security (WS-Security) specification* on their Web sites. This specification included the basic ideas of security token, XML signature, and XML encryption. The specification also defined the format for username tokens and encoded binary security tokens. After some discussion and an interoperability test based on the specification, the following issues were noted:

- The specification requires that the Web services security processors understand the schema correctly so that the processor distinguishes between the ID attribute for XML signature and XML encryption.
- The freshness of the message, which indicates whether the message complies with predefined time constraints, cannot be determined.
- Digested password strings do not strengthen security.

In August 2002, IBM, Microsoft, and VeriSign published the *Web Services Security Addendum*, which attempted to address the previously listed issues. The following solutions were put in the addendum:

- Require a global ID attribute for XML signature and XML encryption
- Use time stamp header elements that indicate the time of the creation, receipt, or expiration of the message
- Use password strings that are digested with a time stamp and nonce (randomly generated token)

### OASIS activities

In June 2002, the Organization for the Advancement of Structured Information Standards (OASIS) received a proposed Web services security specification from IBM, Microsoft, and Verisign. The Web Services Security Technical Committee (WSS TC) was organized at OASIS soon after the submission. The technical committee included many companies including IBM, Microsoft, VeriSign, Sun Microsystems, and BEA Systems.

In September 2002, WSS TC published its first specification, *Web Services Security Core Specification, Working Draft 01*. This specification included the contents of both the original Web services security specification and its addendum.

The coverage of the technical committee became larger as the discussion proceeded. Since the Web Services Security Core Specification allows arbitrary types of security tokens, proposals were published as profiles. The profiles described the method for embedding tokens, including Security Assertion Markup Language (SAML) tokens and Kerberos tokens imbedded into the Web services security messages. Subsequently, the definitions of the usage for user name tokens and X.509 binary security tokens, which were defined in the original Web Services Security Specification, were divided into the profiles.

WebSphere Application Server supports the following specifications:

- Web Services Security: SOAP Message Security Draft 13 (formerly Web Services Security Core Specification)
- Web Services Security: Username Token Profile Draft 2

The following figure shows the various Web services security-related specifications. As indicated in the figure, the current support level for Web services security: SOAP message security is based on Draft 13 from May 2003. The current support level for Web services security User name token profiles, is based on Draft 2 from February 2003.

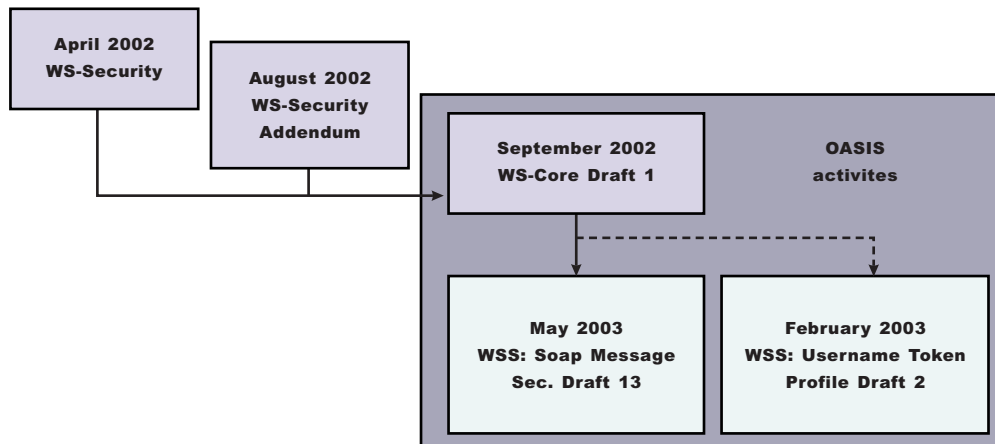


Figure 11. Web services security specification support

## Web services security support

WebSphere Application Server, Versions 4.x, 5, and 5.0.1 support digital signature for Apache Simple Object Access Protocol (SOAP) Version 2.x. Beginning with WebSphere Application Server, Version 5.0.2, IBM supports *Web services security*, which is an extension of the IBM Web services engine to provide a quality of service. The IBM implementation is based on the Web services security specification, "Web Services Security (WS-Security)", originally proposed by IBM, Microsoft, and VeriSign in April 2002. Early versions of the proposed draft specification can be found in Web Services Security (WS-Security) Version 1.0 05 April 2002 and Web Services Security Addendum 18 August 2002. The WebSphere Application Server implementation is based on the Organization for the Advancement of Structured Information Standards (OASIS) working Draft 13 specification. (See the OASIS Web Services Security TC Web site for the latest working specification.) However, not all the features in the OASIS working Draft 13 specification are implemented.

WebSphere Application Server security infrastructure fully integrates Web services security with Java 2 Platform, Enterprise Edition (J2EE) security. When a user ID and password are embedded in a request

message, authentication is performed with the user ID and password. If authentication is successful, a user identity is established and further resource access is authorized based on that identity. After the user ID and password are authenticated by the Web services security run time, a J2EE container performs authorization.

WebSphere Application Server provides an implementation of the key features of Web services security based on the following specifications:

- Specification: Web Services Security (WS-Security) Version 1.0 05 April 2002
- Web Services Security Addendum 18 August 2002
- Web Services Security: SOAP Message Security Working 13 May 2003
- Web Services Security: Username Token Profile Draft

The following table provides a summary of Web services security elements supported by WebSphere Application Server:

*Table 5. Web services security elements*

Element	Notes
UsernameToken	Both the user name and password for the BasicAuth authentication method and the user name for the identity assertion authentication method are supported. WebSphere Application Server, Version 5.1 supports nonce, a randomly generated value.
BinarySecurityToken	X.509 certificates and Lightweight Third Party Authentication (LTPA) can be embedded, but there is no implementation to embed Kerberos tickets. However, the binary token generation and validation are pluggable and are based on the Java Authentication and Authorization Service (JAAS) Application Programming Interfaces (APIs). You can extend this implementation to generate and validate other types of binary security tokens.
Signature	The X.509 certificate is embedded as a binary security token and can be referenced by the SecurityTokenReference. WebSphere Application Server does not support shared, key-based signature.
Encryption	Both the EncryptedKey and ReferenceList XML tags are supported. KeyIdentifier specifies public keys and KeyName identifies the secret keys. WebSphere Application Server has the capability to map an authenticated identity to a key for encryption or use the signer certificate to encrypt the response message.
Timestamp	WebSphere Application Server supports the Created and Expires attributes. The freshness of the message, which indicates whether the message complies with predefined time constraints, is checked only if the Expires attribute is present in the message. WebSphere Application Server does not support the Received attribute, which is defined in the addendum. Instead, WebSphere Application Server uses the TimestampTrace Received attribute, which is defined in the OASIS specification.
XML based token	You can insert and validate an arbitrary format of XML tokens into a message. This format mechanism is based on the JAAS APIs.

Signing and encrypting attachments is not supported by WebSphere Application Server. However, WebSphere Application Server signs and encrypts the following elements for the request message.

Method	Element
XML digital signature	<ul style="list-style-type: none"> <li>• Body</li> <li>• Securitytoken</li> <li>• Timestamp</li> </ul>
XML encryption	<ul style="list-style-type: none"> <li>• Bodycontent</li> <li>• Usenametoken</li> </ul>
AuthMethod	<ul style="list-style-type: none"> <li>• BasicAuth</li> <li>• IDAssertion (From WebSphere Application Server to another WebSphere Application Server)</li> <li>• Signature</li> <li>• Lightweight Third Party Authentication (LTPA) on the server side</li> <li>• Other customer tokens</li> </ul>

WebSphere Application Server signs and encrypts the following elements for the response message:

Method	Element
XML digital signature	<ul style="list-style-type: none"> <li>• Body</li> <li>• Timestamp</li> </ul>
XML encryption	<ul style="list-style-type: none"> <li>• Bodycontent</li> </ul>

The namespaces used for sending a message were published by OASIS in draft 13 (<http://schemas.xmlsoap.org/ws/2003/06/secext>). However, the Web services security run time in WebSphere Application Server can accept any of the following namespaces:

**April 2002 specification**

<http://schemas.xmlsoap.org/ws/2002/04/secext>

**August 2002 addendum**

<http://schemas.xmlsoap.org/ws/2002/07/secext>

<http://schemas.xmlsoap.org/ws/2002/07/utility>

**OASIS draft published on draft 13 May 2003**

<http://schemas.xmlsoap.org/ws/2003/06/secext>

<http://schemas.xmlsoap.org/ws/2003/06/utility>

**Note:** **5.1+** WebSphere Application Server only uses the previously mentioned two name spaces for sending out requests and responses. However, the product can process all other mentioned name spaces for incoming requests and responses.

WebSphere Application Server provides the following capabilities for Web services security:

- Integrity of the message
- Authenticity of the message
- Confidentiality of the message
- Privacy of the message
- Transport level security: provided by Secure Sockets Layer (SSL)
- Security token propagation (pluggable)
- Identity assertion

See the previous table titled, "Web services security elements," for a description of capabilities that are not supported.

## Web services security and Java 2 Platform, Enterprise Edition security relationship

This document describes the relationship between Web services security (message level security) and Java 2 Platform, Enterprise Edition (J2EE) platform security.

WebSphere Application Server supports Java Specification Requests (JSR) 101 and JSR 109 (see Developing Web services for more information). These JSRs define Web services for the Java 2 Platform, Enterprise Edition (J2EE) architecture so that you can develop and run Web services on the J2EE component architecture. "Web services security" refers to the "Web services security: SOAP Message Security" specification (see Web services security support for more information).

### Important

**Note:** "Web services security" refers to the "Web services security: SOAP Message Security" specification (see Web services security support for more information).

### Securing Web services with WebSphere Application Server security (J2EE role-based security)

You can secure Web services using the existing security infrastructure of WebSphere Application Server, J2EE role-based security, and Secure Sockets Layer (SSL) transport level security.

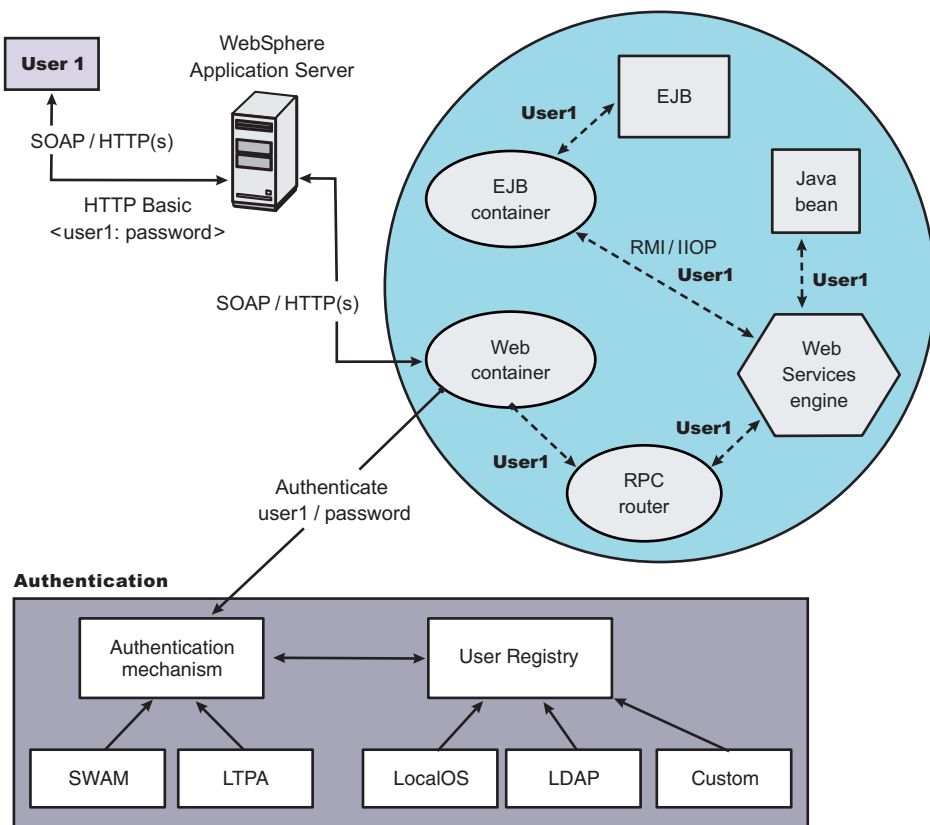


Figure 12. Simple object access protocol message flow using existing security infrastructure of WebSphere Application Server



The Web services endpoint can be secured using J2EE role-based security. The Web services sender sends the basic authentication data using the HTTP header. SSL (HTTPS) can be used to secure the transport. When the WebSphere Application Server receives the SOAP message, the Web container authenticates the user (in this example, user1) and sets the security context for the call. After the security context is set, the SOAP router servlet sends the request to the implementation of the Web services (the implementation can be JavaBeans or enterprise bean files). For enterprise bean implementations, the EJB container performs an authorization check against the identity of user1.

The Web services endpoint also can be secured using the J2EE role. Then, authorization is performed by the Web container before the SOAP request is dispatched to the Web services implementation.

### Securing Web services with Web services security at the message level

You can also secure Web services using Web services security at the message level. In this case, you can digitally sign or encrypt a certain part of the message. Web services security also supports security token propagation within the SOAP message. The following scenario assumes that the Web services endpoint is not secured with J2EE role-based security and the enterprise bean is secured with J2EE role-based security.

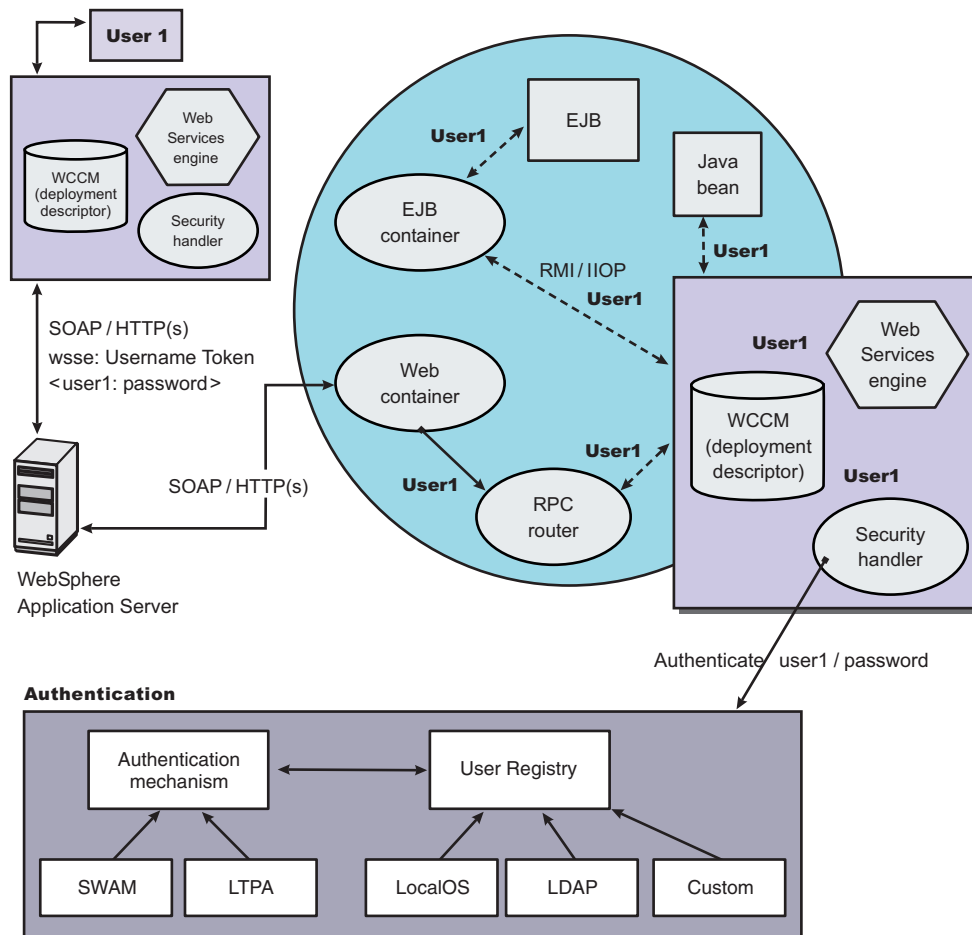


Figure 13. Simple Object Access Protocol message flow using Web services security

In this case, the Web services endpoint is not secured with J2EE role-based security. The Web services engine processes the SOAP message before the client sends the message to the Web services endpoint.

The Web services security run time acts on the security constraints, such as digitally signing, encrypting, or generating (and inserting) a security token in the SOAP header. In this case `<wsse:UsernameToken>` is generated using `user1` and the password. On the server-side (receiving), the Web services process the incoming message and Web services security enforces security constraints. This enforcement includes making sure that messages are properly signed, properly encrypted, and decrypted, authenticating the security token, and setting up the security context with the authenticated identity. (In this case, `user1` is the authenticated identity.) Finally, the SOAP message is dispatched to the Web services implementation (if the implementation is an enterprise beans file, the EJB container performs an authorization check against `user1`). SSL also might be used in this scenario.

### **Mixing the two**

The second scenario shows that Web services security can complement J2EE role-based security. For example, SSL can be enabled at the transport level to provide a secure channel. Furthermore, if the Web services implementation is an enterprise beans file, you can leverage the EJB role-based authorization by performing authorization checks. Web services security run time leverages the security infrastructure to set the authenticated identity in the security context. The authenticated identity can be used in the downstream call to J2EE resources (or other resource types).

There are subtle consequences of combining the two scenarios. For example, if the HTTP transport is sending basic authentication data with `user1` and password in the HTTP header, but `<wsse:UsernameToken>` with `user99` and `letmein` also is inserted into the SOAP header. In the previous scenarios, there are two authentications performed. One authentication is performed by the Web container for authenticating `user1`, and the other is performed by Web services security for authenticating `user99`. The Web services security run time runs after the Web container runs and `user99` is the authenticated identity that is set in the security context.

Web services security can also propagate security tokens from the sender to the receiver for SOAP over a Java Message Service (JMS) transport.

## **Web services security model in WebSphere Application Server**

The Web services security model used by WebSphere Application Server is the declarative model. WebSphere Application Server does not include any application programming interfaces (APIs) for programmatically interacting with Web services security. However, a few Server Provider Interfaces (SPIs) are available for extending some security-related behaviors.

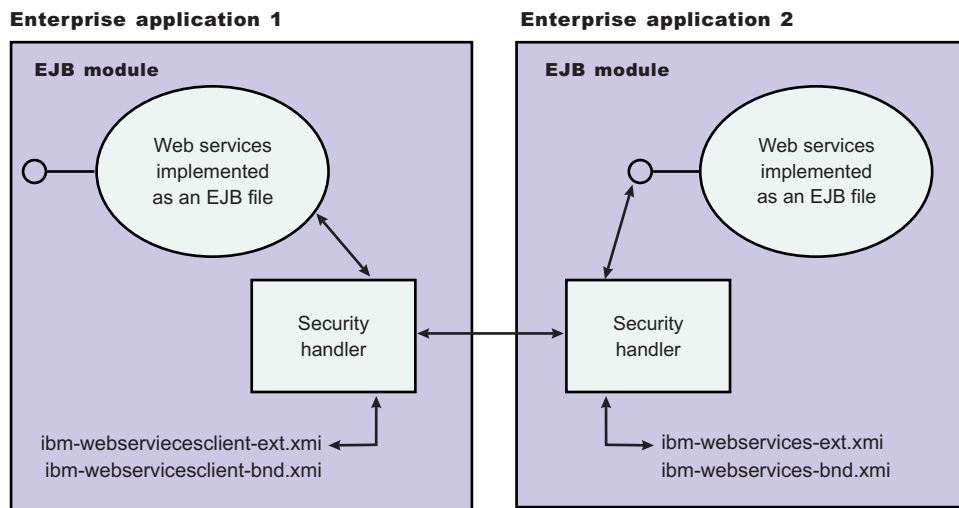


Figure 14. Web services security model

The security constraints for Web services security are specified in IBM deployment descriptor extensions for Web services. The Web services security run time acts on the constraints to enforce Web services security for the Simple Object Access Protocol (SOAP) message. The scope of the IBM deployment descriptor extension is at the enterprise bean (EJB) or Web module level. Bindings are associated with each of the following IBM deployment descriptor extensions:

**Client (Might be either a J2EE Client (Application Client Container) or Web services acting as a client)**

`ibm-webservicesclient-ext.xmi`  
`ibm-webservicesclient-bnd.xmi`

**Server**

`ibm-webservices-ext.xmi`  
`ibm-webservices-bnd.xmi`

It is recommended that you use the tools provided by IBM (the Assembly Toolkit and WebSphere Studio Application Developer) to create the IBM deployment descriptor extension and bindings. After the bindings are created, you can use the administrative console, the Assembly Toolkit, or the WebSphere Studio Application Developer to specify the bindings.

**Important**

**Note:** The binding information is collected after application deployment rather than during application deployment. The alternative is to specify the required binding information before deploying your application.

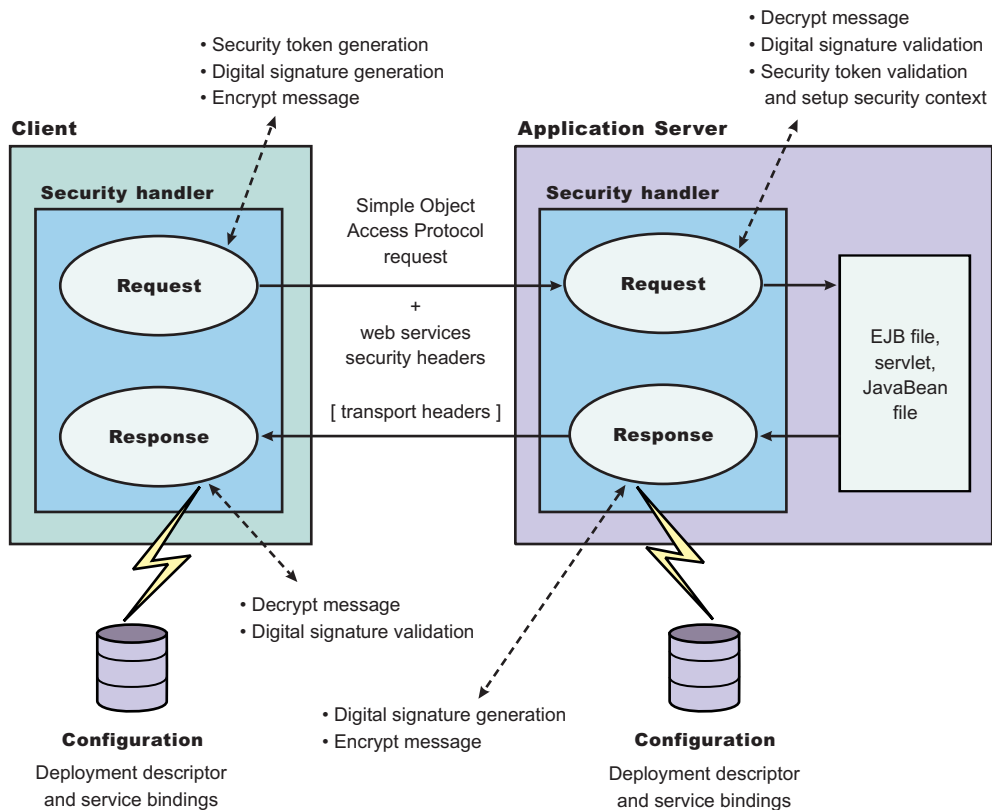


Figure 15. Web services security message interpretation

The Web services security run time enforces Web services security based on the defined security constraints in the deployment descriptor and binding files. Web services security has the following four points where it intercepts the message and acts on the security constraints defined:

Message points	Description
Request sender (defined in the <code>ibm-webservicesclient-ext.xmi</code> and <code>ibm-webservicesclient-bnd.xmi</code> files)	<ul style="list-style-type: none"> <li>Applies the appropriate security constraints to the SOAP message (such as signing or encryption) before the message is sent, generating the time stamp or the required security token.</li> </ul>
Request receiver (defined in the <code>ibm-webservices-ext.xmi</code> and <code>ibm-webservices-bnd.xmi</code> files)	<ul style="list-style-type: none"> <li>Verifies that the Web services security constraints are met.</li> <li>Verifies the freshness of the message based on the time stamp. The freshness of the message indicates whether the message complies with predefined time constraints.</li> <li>Verifies the required signature.</li> <li>Verifies that the message is encrypted and decrypts the message if encrypted.</li> <li>Validates the security tokens and sets up the security context for the downstream call.</li> </ul>
Response sender (defined in the <code>ibm-webservices-ext.xmi</code> and <code>ibm-webservices-bnd.xmi</code> files)	<ul style="list-style-type: none"> <li>Applies the appropriate security constraints to the SOAP message response, like signing the message, encrypting the message, or generating the time stamp.</li> </ul>

Message points	Description
Response receiver (defined in the <code>ibm-webservicesclient-ext.xmi</code> or <code>ibm-webservicesclient-bnd.xmi</code> files)	<ul style="list-style-type: none"> <li>• Verifies that the Web services security constraints are met.</li> <li>• Verifies the freshness of the message based on the time stamp. The freshness of the message indicates whether the message complies with predefined time constraints.</li> <li>• Verifies the required signature.</li> <li>• Verifies that the message is encrypted and decrypts the message, if encrypted.</li> </ul>

## Web services security property collection

Use this page to view a list of additional properties for the configuration.

There are several ways to view a Web services security property collection panel. Complete the following steps to view one of these administrative console pages:

1. Click **Servers > Application Servers > *server\_name***.
2. Under Additional Properties, click **Web Services: Default bindings for Web Services Security > Key Locators > *key\_locator\_name***.
3. Under Additional Properties, click **Properties**.
4. Click **New** to create a new property.
5. Click **Delete** to delete a property that you specified previously.

### Name

Specifies the name of the property.

### Value

Specifies the value for the property.

## Web services security property configuration settings

Use this page to configure additional properties.

There are several ways to view a Web services security property configuration settings panel. Complete the following steps to view one of these administrative console pages:

1. Click **Servers > Application Servers > *server\_name***.
2. Under Additional Properties, click **Web Services: Default bindings for Web Services Security > Key Locators > *key\_locator\_name***.
3. Under Additional Properties, click **Properties > New**.

### Property Name

Specifies the name of the property.

**Data type:** String

### Property Value

Specifies the value for the property.

**Data type:** String

# Usage scenario for propagating security tokens

## A sample scenario

This document describes a usage scenario for Web services security.

In scenario 1, Client 1 invokes Web services 1. Then Web services 1 calls EJB file 2. EJB file 2 calls Web services 3 and Web services 3 calls Web services 4.

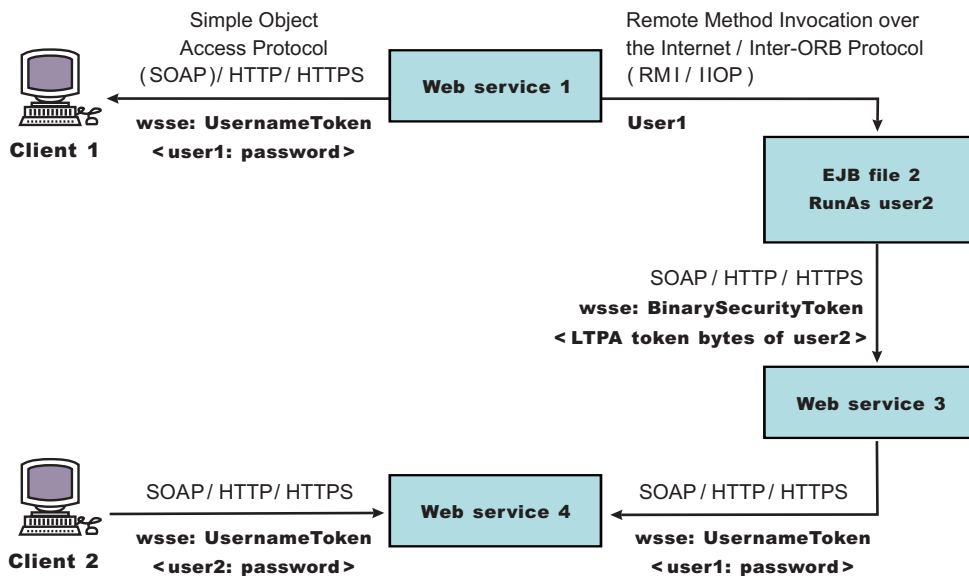


Figure 16. Propagating security tokens

The previous scenario shows how to propagate security tokens using Web services security, the security infrastructure of the WebSphere Application Server, and Java 2 Platform, Enterprise Edition (J2EE) security. Web services 1 is configured to accept `<wsse:UsernameToken>` only and use the BasicAuth authentication method. However, Web services 4 is configured to accept either `<wsse:UsernameToken>` using the BasicAuth authentication method or Lightweight Third Party Authentication (LTPA) as `<wsse:BinarySecurityToken>`. The following steps describe the scenario shown in the previous figure:

1. Client 1 sends a SOAP message to Web services 1 with user1 and password in the `<wsse:UsernameToken>` element.
2. The user1 and password values are authenticated by the Web services security run time and set in the current security context as the Java Authentication and Authorization Service (JAAS) Subject.
3. Web services 1 invokes EJB file 2 using the Remote Method Invocation over the Internet Inter-ORB Protocol (RMI/IIOP) protocol.
4. The user1 identity is propagated to the downstream call.
5. The EJB container of EJB file 2 performs an authorization check against user1.
6. EJB file 2 calls Web services 3 and Web services 3 is configured to accept LTPA tokens.
7. The RunAs role of EJB file 2 is set to user2.
8. The LTPA CallbackHandler implementation extracts the LTPA token from the current JAAS Subject in the security context and Web services security run time inserts the token as `<wsse: BinarySecurityToken>` in the SOAP header.
9. The Web services security run time in Web services 3 calls the JAAS login configuration to validate the LTPA token and set it in the current security context as the JAAS Subject.

10. Web services 3 is configured to send LTPA security to Web services 4. In this case, assume that the RunAs role is not configured for Web services 3. The LTPA token of user2 is propagated to Web services 4.
11. Client 2 uses the <wsse:UsernameToken> element to propagate the basic authentication data to Web services 4.

Web services security complements the WebSphere Application Server security run time and the J2EE role-based security. This scenario demonstrates how to propagate security tokens across multiple resources such as Web services and EJB files.

## Configurations

The Web services security model used by WebSphere Application Server is the declarative model.

No Application Programming Interfaces (APIs) exist in WebSphere Application Server for programmatically interacting with Web services security. However, Service Provider Programming Interfaces (SPIs) are available for extending some security run-time behaviors. You can secure an application with Web services security by defining security constraints in the IBM extension deployment descriptors and in IBM extension bindings.

The development life cycle of a Web services security-enabled application is similar to the Java 2 Platform, Enterprise Edition (J2EE) model. See the following figure for more details.

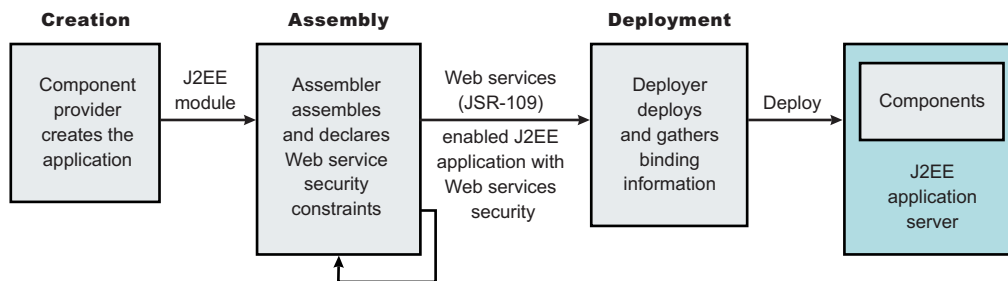


Figure 17. Application development life cycle

The Web services security constraints are defined by the assembler during the application assembly phase if the J2EE application is Web services-enabled. Create, define, and edit the Web services security constraints with the Assembly Toolkit, which can be downloaded from the following location:[http://www-1.ibm.com/support/docview.wss?rs=180&context=SSEQTP&q=ASTK&uid=swg24005125&loc=en\\_US&cs=utf-8&lang=en+en](http://www-1.ibm.com/support/docview.wss?rs=180&context=SSEQTP&q=ASTK&uid=swg24005125&loc=en_US&cs=utf-8&lang=en+en)

### Web services security constraints

The security constraints for Web services security are specified in the IBM deployment descriptor extension for Web services. The assembler defines these constraints during the application assembly phase, if the J2EE application is Web services enabled. Define the Web services security constraints using the Assembly Toolkit.

The Web services security run time acts on the constraints to enforce Web services security for the SOAP message. The scope of the IBM deployment descriptor extension is at the EJB module or Web module level. There also are bindings associated with each of the following IBM deployment descriptor extensions:

**Client** (might be either a J2EE client (application client container) or Web services acting as a client)



- ibm-webservicesclient-ext.xmi
- ibm-webservicesclient-bnd.xmi

**Server**

- ibm-webservices-ext.xmi
- ibm-webservices-bnd.xmi

The IBM extension deployment descriptor and bindings are associated with each EJB module or Web module. See Figure 2 for more information. If Web services is acting as a client, then it contains the client IBM extension deployment descriptors and bindings in the EJB module or Web module.

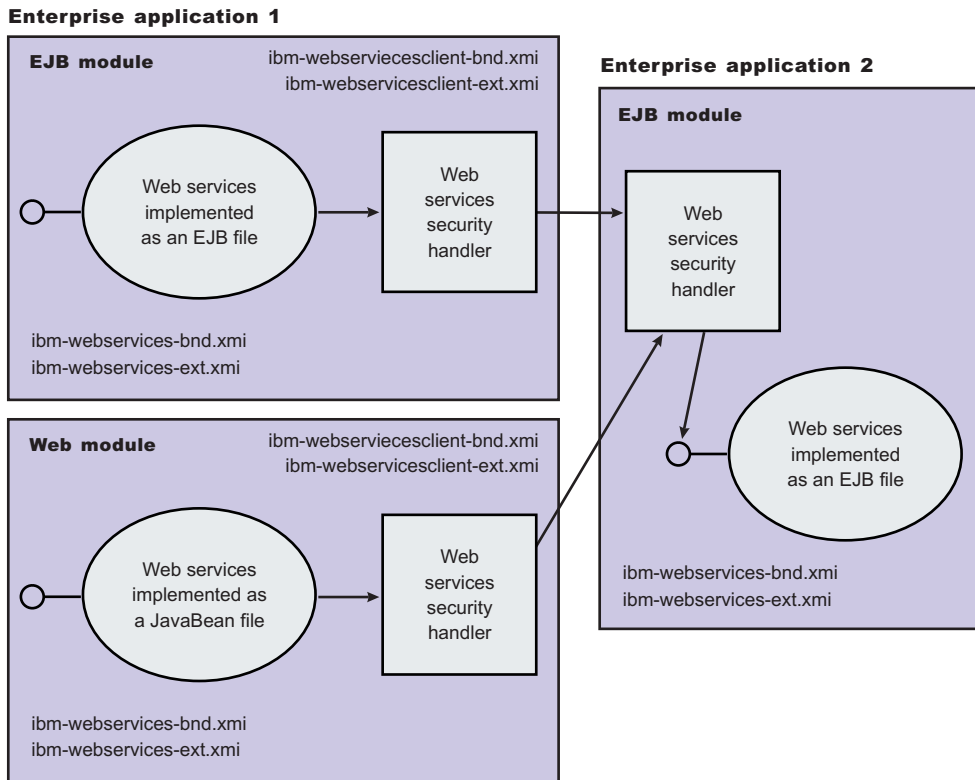


Figure 18. IBM extension deployment descriptors and bindings

The Web services security handler acts on the security constraints defined in the IBM extension deployment descriptor and enforces the security constraints accordingly. There are outbound and inbound configurations in both the client and server security constraints.

In a SOAP request, the following message points exist:

- Sender outbound
- Receiver inbound
- Receiver outbound
- Sender inbound

These message points correspond to the following four security constraints:

- Request sender (sender outbound)
- Request receiver (receiver inbound)

- Response sender (receiver outbound)
- Response receiver (sender inbound)

The security constraints of request sender and request receiver must match. Also, the security constraints of the response sender and response receiver must match. For example, if you specify integrity as a constraint in the request receiver, then you must configure the request sender to have integrity applied to the SOAP message. Otherwise, the request is denied because the SOAP message does not include the integrity specified in the request constraint.

The four security constraints are shown in the following figure of Web services security constraints.

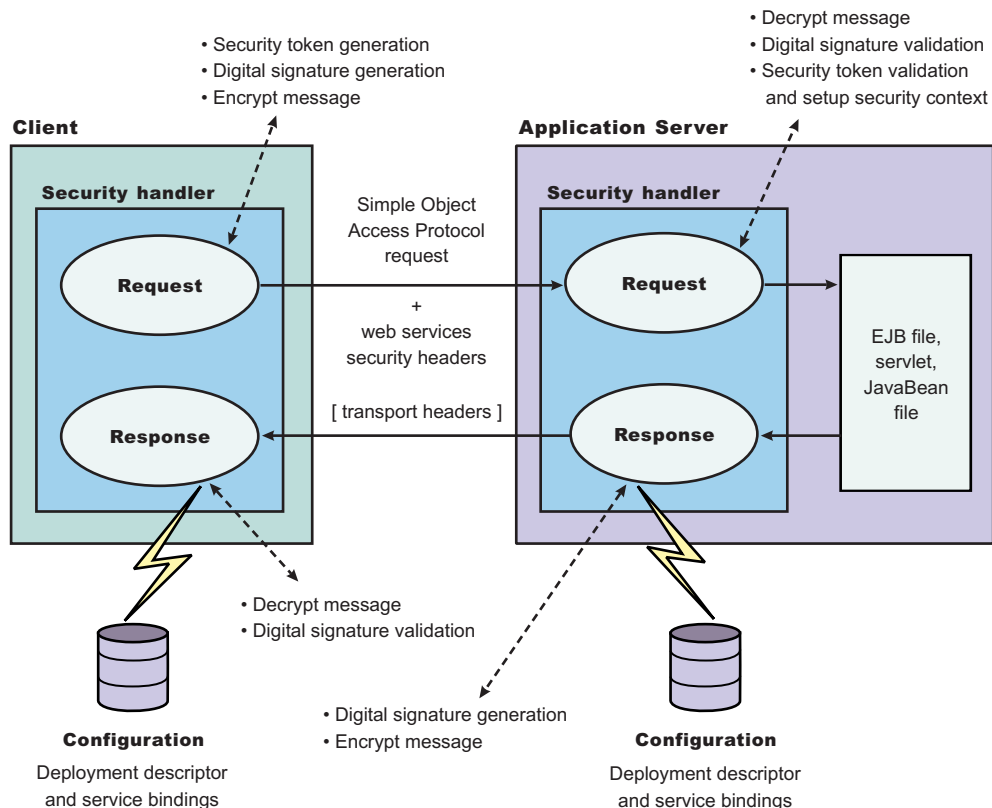


Figure 19. Web services security constraints

## Sample configuration

WebSphere Application Server provides the following sample key stores for sample configurations. These sample key stores are for testing and sample purposes only. Do not use them in a production environment.

- {USER\_INSTALL\_ROOT}/etc/ws-security/samples/dsig-sender.ks
  - The keystore password is client
  - Trusted certificate with alias name, soapca
  - Personal certificate with alias name, soaprequester and key password client issued by intermediary certificate authority Int CA2, which is, in turn, issued by soapca
- {USER\_INSTALL\_ROOT}/etc/ws-security/samples/dsig-receiver.ks
  - The keystore password is server
  - Trusted certificate with alias name, soapca

- Personal certificate with alias name, soapprovider and key password server, issued by intermediary certificate authority Int CA2, which is, in turn, issued by soapca
- {USER\_INSTALL\_ROOT}/etc/ws-security/samples/enc-sender.jceks
  - The keystore password is storepass
  - Secret key CN=Group1, alias name Group1, and key password keypass
  - Public key CN=Bob, O=IBM, C=US, alias name bob, and key password keypass
  - Private key CN=Alice, O=IBM, C=US, alias name alice, and key password keypass
- {USER\_INSTALL\_ROOT}/etc/ws-security/samples/enc-receiver.jceks
  - The keystore password is storepass
  - Secret key CN=Group1, alias name Group1, and key password keypass
  - Private key CN=Bob, O=IBM, C=US, alias name bob, and key password keypass
  - Public key CN=Alice, O=IBM, C=US, alias name alice, and key password keypass
- {USER\_INSTALL\_ROOT}/etc/ws-security/samples/intca2.cer
  - The intermediary certificate authority is Int CA2.

### Default binding (cell and server level)

WebSphere Application Server provides the following default binding information:

#### Trust anchors

Used to validate the trust of the signer certificate.

- SampleClientTrustAnchor is used by the response receiver to validate the signer certificate.
- SampleServerTrustAnchor is used by the request receiver to validate the signer certificate.

#### Collection Certificate Store

Used to validate the certificate path.

- SampleCollectionCertStore is used by the response receiver and the request receiver to validate the signer certificate path.

#### Key Locators

Used to locate the key for signature, encryption, and decryption.

- SampleClientSignerKey is used by the requesting sender to sign the SOAP message. The signing key name is clientsignerkey, which can be referenced in the signing information as the signing key name.
- SampleServerSignerKey is used by the responding sender to sign the SOAP message. The signing key name is serversignerkey, which can be referenced in the signing information as the signing key name.
- SampleSenderEncryptionKeyLocator is used by the sender to encrypt the SOAP message. It is configured to use the {USER\_INSTALL\_ROOT}/etc/ws-security/samples/enc-sender.jceks keystore and the com.ibm.wsspi.wssecurity.config.KeyStoreKeyLocator keystore key locator.
- SampleReceiverEncryptionKeyLocator is used by the receiver to decrypt the encrypted SOAP message. The implementation is configured to use the {USER\_INSTALL\_ROOT}/etc/ws-security/samples/enc-receiver.jceks keystore and the com.ibm.wsspi.wssecurity.config.KeyStoreKeyLocator keystore key locator. The implementation is configured for symmetric encryption (DES or TRIPLEDES). However, to use it for asymmetric encryption (RSA), you must add the private key CN=Bob, O=IBM, C=US, alias name bob, and key password keypass.
- SampleResponseSenderEncryptionKeyLocator is used by the response sender to encrypt the SOAP response message. It is configured to use the {USER\_INSTALL\_ROOT}/etc/ws-security/samples/enc-receiver.jceks keystore and the com.ibm.wsspi.wssecurity.config.WSIdKeyStoreMapKeyLocator key locator. This key locator maps an authenticated identity (of the current thread) to a public key for encryption. By default,

WebSphere Application Server is configured to map to public key `alice`, and you must change WebSphere Application Server to the appropriate user. The `SampleResponseSenderEncryptionKeyLocator` key locator also can set a default key for encryption. By default, this key locator is configured to use public key `alice`.

### Trusted ID Evaluator

Used to establish trust before asserting to the identity in identity assertion.

`SampleTrustedIDEvaluator` is configured to use the `com.ibm.wsspi.wssecurity.id.TrustedIDEvaluatorImpl` implementation. The default implementation of `com.ibm.wsspi.wssecurity.id.TrustedIDEvaluator` contains a list of trusted identities. The list is defined as properties with `trustedId_*` as the key and the value as the trusted identity. Define this information for the server level in the administration console by completing the following steps:

1. Click **Servers > Application Servers > *server1***.
2. Under Additional Properties, click **Web Services: Default bindings for Web Services Security > Trusted ID Evaluators > *SampleTrustedIDEvaluator***

### Login Mapping

Used to authenticate the incoming security token in the Web services security SOAP header of a SOAP message.

- The `BasicAuth` authentication method is used to authenticate user name security token (user name and password).
- The signature authentication method is used to map a distinguished name (DN) into a WebSphere Application Server Java Authentication and Authorization Server (JAAS) Subject.
- The `IDAssertion` authentication method is used to map a trusted identity into a WebSphere Application Server JAAS Subject for identity assertion.
- The Lightweight Third Party Authentication (LTPA) authentication method is used to validate a LTPA security token.

The previous default bindings for trust anchors, collection certificate stores, and key locators are for testing or sample purpose only. Do not use them for production.

### A sample configuration

The following examples demonstrate what IBM deployment descriptor extensions and bindings can do. The unnecessary information was removed from the examples to improve clarity. Do not copy and paste these examples into your application deployment descriptors or bindings. These examples serve as reference only and are not representative of the recommended configuration.

Use the following tools to create or edit IBM deployment descriptor extensions and bindings:

- Use the Assembly Toolkit to create or edit the IBM deployment descriptor extensions.
- Use the Assembly Toolkit or the administrative console to create or edit the bindings file.

The following example illustrates a scenario that:

- Signs the SOAP body, time stamp, and security token.
- Encrypts the body content and user name token.
- Sends the user name token (basic authentication data).
- Generates the time stamp for the request.

For the response, the SOAP body and time stamp are signed, the body content is encrypted, and the SOAP message freshness is checked using the time stamp. The freshness of the message indicates whether the message complies with predefined time constraints.

The request sender and the request receiver are a pair. Similarly, the response sender and the response receiver are a pair.

**Tip:** It is recommended that you use the WebSphere Application Server variables for specifying the path to the key stores. In the administrative console, click **Environment > Manage WebSphere Variables**. These variables often help with platform differences such as file system naming conventions. In the following examples, `${USER_INSTALL_ROOT}` is used for specifying the path to the key stores.

### Client-side IBM deployment descriptor extension

The client-side IBM deployment descriptor extension describes the following constraints:

#### Request Sender

- Signs the SOAP body, time stamp and security token
- Encrypts the body content and user name token
- Sends the basic authentication token (user name and password)
- Generates the time stamp to expire in three minutes

#### Response Receiver

- Verifies that the SOAP body and time stamp are signed
- Verifies that the SOAP body content is encrypted
- Verifies that the time stamp is present (also check for message freshness)

### Example 1: Sample client IBM deployment descriptor extension

The `xmi:id` statements are removed for readability. These statements must be added for this example to work.

**Important:** In the following code sample, lines 2 through 4 were split into three lines due to the width of the printed page.

```
<?xml version="1.0" encoding="UTF-8"?>
<com.ibm.etools.webservice.wssect:WsClientExtension xmi:version="2.0"
 xmlns:xmi="http://www.omg.org/XMI" xmlns:com.ibm.etools.webservice.wssect=
 http://www.ibm.com/websphere/appserver/schemas/5.0.2/wssect.xmi">
 <serviceRefs serviceRefLink="service/myServ">
 <portQnameBindings portQnameLocalNameLink="Port1">
 <clientServiceConfig actorURI="myActorURI">
 <securityRequestSenderServiceConfig actor="myActorURI">
 <integrity>
 <references part="body"/>
 <references part="timestamp"/>
 <references part="securitytoken"/>
 </integrity>
 <confidentiality>
 <confidentialParts part="bodycontent"/>
 <confidentialParts part="usernameToken"/>
 </confidentiality>
 <loginConfig authMethod="BasicAuth"/>
 <addCreatedTimeStamp flag="true" expires="PT3M"/>
 </securityRequestSenderServiceConfig>
 <securityResponseReceiverServiceConfig>
 <requiredIntegrity>
 <references part="body"/>
 <references part="timestamp"/>
 </requiredIntegrity>
 <requiredConfidentiality>
 <confidentialParts part="bodycontent"/>
 </requiredConfidentiality>
 </securityResponseReceiverServiceConfig>
 </clientServiceConfig>
 </portQnameBindings>
 </serviceRefs>
</com.ibm.etools.webservice.wssect:WsClientExtension>
```

```

 <addReceivedTimeStamp flag="true"/>
 </securityResponseReceiverServiceConfig>
</clientServiceConfig>
</portQnameBindings>
</serviceRefs>
</com.ibm.etools.webservice.wsclient:WsClientExtension>

```

## Client-side IBM extension bindings

Example 2 shows the client-side IBM extension binding for the security constraints described previously in the discussion on client-side IBM deployment descriptor extensions.

The signer key and encryption (decryption) key for the message can be obtained from the keystore key locator implementation (`com.ibm.wsspi.wssecurity.config.KeyStoreKeyLocator`). The signer key is used for encrypting the response. The sample is configured to use the Java Certification Path API to validate the certificate path of the signer of the digital signature. The user name token (basic authentication) data is collected from the standard in (stdin) prompts using one of the default Java Authentication and Authorization Service (JAAS) implementations: `javax.security.auth.callback.CallbackHandler` implementation (`com.ibm.wsspi.wssecurity.auth.callback.StdinPromptCallbackHandler`).

### Example 2: Sample client IBM extension binding

**Important:** In the following code sample, several lines were split into multiple lines due to the width of the printed page. See the close bracket for an indication of where each line of code ends.

```

<?xml version="1.0" encoding="UTF-8"?>
<com.ibm.etools.webservice.wsclient:ClientBinding xmi:version="2.0"
 xmlns:xmi="http://www.omg.org/XMI"
 xmlns:com.ibm.etools.webservice.wsclient="
 http://www.ibm.com/websphere/appserver/schemas/5.0.2/wsclient.xmi">
 <serviceRefs serviceRefLink="service/MyServ">
 <portQnameBindings portQnameLocalNameLink="Port1">
 <securityRequestSenderBindingConfig>
 <signingInfo>
 <signatureMethod algorithm="http://www.w3.org/2000/09/xmldsig#rsa-sha1"/>
 <signingKey name="clientsignerkey" locatorRef="SampleClientSignerKey"/>
 <canonicalizationMethod algorithm="http://www.w3.org/2001/10/xml-exc-c14n#"/>
 <digestMethod algorithm="http://www.w3.org/2000/09/xmldsig#sha1"/>
 </signingInfo>
 <keyLocators name="SampleClientSignerKey" classname="
 com.ibm.wsspi.wssecurity.config.KeyStoreKeyLocator">
 <keyStore storepass="{xor}PDM20jEr" path="
 ${USER_INSTALL_ROOT}/etc/ws-security/samples/dsig-sender.ks" type="JKS"/>
 <keys alias="soaprequester" keypass="{xor}PDM20jEr" name="clientsignerkey"/>
 </keyLocators>
 <encryptionInfo name="EncInfo1">
 <encryptionKey name="CN=Bob, O=IBM, C=US" locatorRef="
 SampleSenderEncryptionKeyLocator"/>
 <encryptionMethod algorithm="http://www.w3.org/2001/04/xmlenc#tripledes-cbc"/>
 <keyEncryptionMethod algorithm="http://www.w3.org/2001/04/xmlenc#rsa-1_5"/>
 </encryptionInfo>
 <keyLocators name="SampleSenderEncryptionKeyLocator" classname="
 com.ibm.wsspi.wssecurity.config.KeyStoreKeyLocator">
 <keyStore storepass="{xor}LCswLTovPiws" path="
 ${USER_INSTALL_ROOT}/etc/ws-security/samples/enc-sender.jceks" type="JCEKS"/>
 <keys alias="Group1" keypass="{xor}NDomLz4sLA==" name="CN=Group1"/>
 </keyLocators>
 <loginBinding authMethod="BasicAuth" callbackHandler="
 com.ibm.wsspi.wssecurity.auth.callback.StdinPromptCallbackHandler"/>
 </securityRequestSenderBindingConfig>
 <securityResponseReceiverBindingConfig>
 <signingInfos>
 <signatureMethod algorithm="http://www.w3.org/2000/09/xmldsig#rsa-sha1"/>
 <certPathSettings>

```

```

 <trustAnchorRef ref="SampleClientTrustAnchor"/>
 <certStoreRef ref="SampleCollectionCertStore"/>
 </certPathSettings>
 <canonicalizationMethod algorithm="http://www.w3.org/2001/10/xml-exc-c14n#"/>
 <digestMethod algorithm="http://www.w3.org/2000/09/xmlsig#sha1"/>
</signingInfos>
<trustAnchors name="SampleClientTrustAnchor">
 <keyStore storepass="{xor}PDM20jEr" path=
 "${USER_INSTALL_ROOT}/etc/ws-security/samples/dsig-sender.ks" type="JKS"/>
</trustAnchors>
<certStoreList>
 <collectionCertStores provider="IBMCertPath" name="SampleCollectionCertStore">
 <x509Certificates path="${USER_INSTALL_ROOT}/etc/ws-security/samples/intca2.cer"/>
 </collectionCertStores>
</certStoreList>
<encryptionInfos name="EncInfo2">
 <encryptionKey locatorRef="SampleReceiverEncryptionKeyLocator"/>
 <encryptionMethod algorithm="http://www.w3.org/2001/04/xmlenc#tripledes-cbc"/>
 <keyEncryptionMethod algorithm="http://www.w3.org/2001/04/xmlenc#rsa-1_5"/>
</encryptionInfos>
<keyLocators name="SampleReceiverEncryptionKeyLocator" classname=
 "com.ibm.wsspi.wssecurity.config.KeyStoreKeyLocator">
 <keyStore storepass="{xor}PDM20jEr" path=
 "${USER_INSTALL_ROOT}/etc/ws-security/samples/dsig-sender.ks" type="JKS"/>
 <keys alias="soaprequester" keypass="{xor}PDM20jEr" name="clientsignerkey"/>
</keyLocators>
</securityResponseReceiverBindingConfig>
</portQnameBindings>
</serviceRefs>
</com.ibm.etools.webservice.wscbind:ClientBinding>

```

## Server-side IBM deployment descriptor extension

The client-side IBM deployment descriptor extension describes the following constraints:

Request Receiver (ibm-webservices-ext.xmi and ibm-webservices-bnd.xmi)

- Verifies that the SOAP body, time stamp, and security token are signed.
- Verifies that the SOAP body content and user name token are encrypted.
- Verifies that the basic authentication token (user name and password) is in the Web services security SOAP header.
- Verifies that the time stamp is present (also check for message freshness). The freshness of the message indicates whether the message complies with predefined time constraints.

Response Sender (ibm-webservices-ext.xmi and ibm-webservices-bnd.xmi)

- Signs the SOAP body and time stamp
- Encrypts the SOAP body content
- Generates the time stamp to expire in 3 minutes

### Example 3: Sample server IBM deployment descriptor extension

**Important:** In the following code sample, several lines were split into multiple lines due to the width of the printed page. See the close bracket for an indication of where each line of code ends.

```

<?xml version="1.0" encoding="UTF-8"?>
<com.ibm.etools.webservice.wsext:WsExtension xmi:version="2.0"
 xmlns:xmi="http://www.omg.org/XMI"
 xmlns:com.ibm.etools.webservice.wsext=
http://www.ibm.com/websphere/appserver/schemas/5.0.2/wsext.xmi">
 <wsDescExt wsDescNameLink="MyServ">
 <pcBinding pcNameLink="Port1">
 <serverServiceConfig actorURI="myActorURI">

```



```

<securityRequestReceiverServiceConfig>
 <requiredIntegrity>
 <references part="body"/>
 <references part="timestamp"/>
 <references part="securitytoken"/>
 </requiredIntegrity>
 <requiredConfidentiality">
 <confidentialParts part="bodycontent"/>
 <confidentialParts part="usernameToken"/>
 </requiredConfidentiality>
 <loginConfig>
 <authMethods text="BasicAuth"/>
 </loginConfig>
 <addReceivedTimestamp flag="true"/>
</securityRequestReceiverServiceConfig>
<securityResponseSenderServiceConfig actor="myActorURI">
 <integrity>
 <references part="body"/>
 <references part="timestamp"/>
 </integrity>
 <confidentiality>
 <confidentialParts part="bodycontent"/>
 </confidentiality>
 <addCreatedTimestamp flag="true" expires="PT3M"/>
</securityResponseSenderServiceConfig>
</serverServiceConfig>
</pcBinding>
</wsDescExt>
</com.ibm.etools.webservice.wsExt:WsExtension>

```

## Server-side IBM extension bindings

The following binding information reuses some of the default binding information defined either at the server level or the cell level, which depends upon the installation. For example, request receiver is referencing the `SampleCollectionCertStore` certification store and the `SampleServerTrustAnchor` trust store is defined in the default binding. However, the encryption information in the request receiver is referencing a `SampleReceiverEncryptionKeyLocator` key locator defined in the application-level binding (the same `ibm-webservices-bnd.xmi` file). The response sender is configured to use the signer key of the digital signature of the request to encrypt the response using one of the default key locator (`com.ibm.wsspi.wssecurity.config.CertInRequestKeyLocator`) implementations.

### Example 4: Sample server IBM extension binding

```

<?xml version="1.0" encoding="UTF-8"?>
<com.ibm.etools.webservice.wsbind:WSBinding xmi:version="2.0"
 xmlns:xmi="http://www.omg.org/XMI"
 xmlns:com.ibm.etools.webservice.wsbind=
 http://www.ibm.com/websphere/appserver/schemas/5.0.2/wsbind.xmi">
 <wsdescBindings wsDescNameLink="MyServ">
 <pcBindings pcNameLink="Port1" scope="Session">
 <securityRequestReceiverBindingConfig>
 <signingInfos>
 <signatureMethod algorithm="http://www.w3.org/2000/09/xmldsig#rsa-sha1"/>
 <certPathSettings>
 <trustAnchorRef ref="SampleServerTrustAnchor"/>
 <certStoreRef ref="SampleCollectionCertStore"/>
 </certPathSettings>
 <canonicalizationMethod algorithm="http://www.w3.org/2001/10/xml-exc-c14n#"/>
 <digestMethod algorithm="http://www.w3.org/2000/09/xmldsig#sha1"/>
 </signingInfos>
 <encryptionInfos name="EncInfo1">
 <encryptionKey locatorRef="SampleReceiverEncryptionKeyLocator"/>
 <encryptionMethod algorithm="http://www.w3.org/2001/04/xmlenc#tripleDES-cbc"/>
 <keyEncryptionMethod algorithm="http://www.w3.org/2001/04/xmlenc#rsa-1_5"/>
 </encryptionInfos>
 </securityRequestReceiverBindingConfig>
 </pcBindings>
 </wsdescBindings>
</com.ibm.etools.webservice.wsbind:WSBinding>

```

```

<keyLocators name="SampleReceiverEncryptionKeyLocator" classname=
"com.ibm.wsspi.wssecurity.config.KeyStoreKeyLocator">
 <keyStore storepass="{xor}LCswLTovPiws" path="{USER_INSTALL_ROOT}/
etc/ws-security/samples/enc-receiver.jceks" type="JCEKS"/>
 <keys alias="Group1" keypass="{xor}NDomLz4sLA==" name="CN=Group1"/>
 <keys alias="bob" keypass="{xor}NDomLz4sLA==" name="CN=Bob, O=IBM, C=US"/>
</keyLocators>
</securityRequestReceiverBindingConfig>
<securityResponseSenderBindingConfig>
 <signingInfo>
 <signatureMethod algorithm="http://www.w3.org/2000/09/xmlsig#rsa-sha1"/>
 <signingKey name="serversignerkey" locatorRef="SampleServerSignerKey"/>
 <canonicalizationMethod algorithm="http://www.w3.org/2001/10/xml-exc-c14n#"/>
 <digestMethod algorithm="http://www.w3.org/2000/09/xmlsig#sha1"/>
 </signingInfo>
 <encryptionInfo name="EncInfo2">
 <encryptionKey locatorRef="SignerKeyLocator"/>
 <encryptionMethod algorithm="http://www.w3.org/2001/04/xmlenc#tripledes-cbc"/>
 <keyEncryptionMethod algorithm="http://www.w3.org/2001/04/xmlenc#rsa-1_5"/>
 </encryptionInfo>
 <keyLocators name="SignerKeyLocator" classname=
"com.ibm.wsspi.wssecurity.config.CertInRequestKeyLocator"/>
</securityResponseSenderBindingConfig>
</pcBindings>
</wsdescBindings>
<routerModules transport="http" name="StockQuote.war"/>
</com.ibm.etools.webservice.wsbind:WSBinding>

```

## View Web services client deployment descriptor

Use this page to view your client deployment descriptor.

To view this administrative console page, complete the following steps:

1. Click **Applications > Enterprise Applications** > *application\_name*.
2. Under Related Items, click **Web Modules > URI\_file\_name > View Web Services Client Deployment Descriptor**.

The information in the following implementation descriptions indicates how to configure your application-level bindings. If the Web server is acting as a client, the default bindings are used. To configure the server-level bindings, which are the defaults, complete the following steps:

1. Click **Servers > Application Servers** > *server\_name*.
2. Under Related Items, click **Web Services: Default bindings for Web Services Security**.

If you are using any of the following configurations, verify that the deployment descriptor is configured properly:

- Request signing
- Request encryption
- BasicAuth authentication
- Identity (ID) Assertion authentication
- Identity (ID) Assertion authentication with the signature TrustMode
- Response digital signature verification
- Response decryption

## Request signing

If the integrity constraints (digital signature) are specified, verify that you configured the signing information in the binding files.

To configure the signing parameters, complete the following steps:

1. Click **Applications > Enterprise Applications** > *application\_name*.

2. Under Related Items, click **Web Modules** >*URI\_file\_name* **Web Services: Client Security Bindings** .
3. In the Response Receiver Binding column, click **Edit** > **Signing Information** > **New**.

To configure the key locators, complete the following steps:

1. Click **Servers** > **Application Servers** > *server\_name*.
2. Under Additional Properties, click **Web Services: Default bindings for Web Services Security** > **Key Locators**.

### Request encryption

If the confidentiality constraints (encryption) are specified, verify that you configured the encryption information in the binding files.

To configure the encryption parameters, complete the following steps:

1. Click **Applications** > **Enterprise Applications** >*application\_name*.
2. Under Related Items, click **Web Modules** >*URI\_file\_name* > **Web Services: Client Security Bindings** .
3. In the Response Receiver Binding column, click **Edit** > **Encryption Information** > **New**.

To configure the key locators, complete the following steps:

1. Click **Servers** > **Application Servers** > *server\_name*.
2. Under Additional Properties, click **Web Services: Default bindings for Web Services Security** > **Key Locators**.

### BasicAuth authentication

If BasicAuth authentication is configured as the required security token, specify the CallbackHandler in the binding file to collect the basic authentication data. The following list contains the Callback support implementations:

#### **com.ibm.wsspi.wssecurity.auth.callback.GuiPromptCallbackHandler**

This implementation prompts for BasicAuth information (user name and password) in an interface.

#### **com.ibm.wsspi.wssecurity.auth.callback.NonPromptCallbackHandler**

This implementation reads the BasicAuth information from the binding file.

#### **com.ibm.wsspi.wssecurity.auth.callback.StdPromptCallbackHandler**

This implementation prompts for a user name and password using the standard in (stdin) prompt.

To configure the login binding information, complete the following steps:

1. Click **Applications** > **Enterprise Applications** >*application\_name*.
2. Under Related Items, click **Web Module** >*URI\_file\_name* > **Web Services: Client Security Bindings**.
3. Under Request Sender Bindings, click **Edit** > **Login Binding**.

### Identity (ID) Assertion authentication with BasicAuth TrustMode

Configure a login binding in the bindings file with a `com.ibm.wsspi.wssecurity.auth.callback.NonPromptCallbackHandler` implementation. Specify a BasicAuth user ID and password that a TrustedIDEvaluator on a downstream server trusts.

To configure the login binding information, complete the following steps:

1. Click **Applications** > **Enterprise Applications** >*application\_name*.
2. Under Related Items, click **Web Module** >*URI\_file\_name*> **Web Services: Client Security Bindings**.
3. Under **Request Sender Bindings**, click **Edit** > **Login Binding**.

## Identity (ID) Assertion authentication with the Signature TrustMode

Configure the signing information in the bindings file with a signing key pointing to a key locator. The key locator contains the X.509 certificate that is trusted by the downstream server.

To configure ID assertion, complete the following steps:

1. Click **Servers > Application Servers > *server\_name***.
2. Under Additional Properties, click **Web Services: Default bindings for Web Services Security > Login Mappings > IDAssertion**.

To configure the login binding information, complete the following steps:

1. Click **Applications > Enterprise Applications > *application\_name***.
2. Under Related Items, click **Web Module > *URI\_file\_name* > Web Services: Client Security Bindings**.
3. Under Request Sender Bindings, click **Edit > Login Binding**.

## Response digital signature verification

If the integrity constraints (signature required) are defined, verify that you configured the signing information in the binding files.

To configure the signing parameters, complete the following steps:

1. Click **Applications > Enterprise Applications > *application\_name***.
2. Under Related Items, click **Web Modules > *URI\_file\_name* > Web Services: Client Security Bindings**.
3. In the Response Receiver Binding column, click **Edit > Signing Information > New**.

To configure the trust anchors, complete the following steps:

1. Click **Servers > Application Servers > *server\_name***.
2. Under Additional Properties, click **Web Services: Default bindings for Web Services Security > Trust Anchors > New**.

To configure the collection certificate store, complete the following steps:

1. Click **Servers > Application Servers > *server\_name***.
2. Under Additional Properties, click **Web Services: Default bindings for Web Services Security > Collection Certificate Store > New**.

## Response decryption

If the confidentiality constraints (encryption) are specified, verify that you defined the encryption information.

To configure the encryption information, complete the following steps:

1. Click **Applications > Enterprise Applications > *application\_name***.
2. Under Related Items, click **Web Modules > *URI\_file\_name* > Web Services: Client Security Bindings**.
3. In the Response Receiver Binding column, click **Edit > Encryption Information > New**.

To configure the key locators, complete the following steps:

1. Click **Servers > Application Servers > *server\_name***.
2. Under Additional Properties, click **Web Services: Default bindings for Web Services Security > Key Locators**.

## View Web services server deployment descriptor

Use this page to view your server deployment descriptor settings.

To view this administrative console page, complete the following steps:

1. Click **Applications > Enterprise Applications > *application\_name***.
2. Under Related Items, click **Web Modules > *URI\_file\_name* > View Web Services Server Deployment Descriptor**.

The information in the following implementation descriptions indicate how to configure your application-level bindings. To configure the server-level bindings, which are the defaults, complete the following steps:

1. Click **Servers > Application Servers > *server\_name***.
2. Under Related Items, click **Web Services: Default bindings for Web Services Security**.
  - Request digital signature verification
  - Request decryption
  - BasicAuth authentication
  - Identity (ID) Assertion authentication
  - Identity (ID) Assertion authentication with the signature TrustMode
  - Response signing
  - Response encryption

### Request digital signature verification

If the integrity constraints (signature required) are defined, verify that you configured the signing information in the binding files.

To configure the signing parameters, complete the following steps:

1. Click **Applications > Enterprise Applications > *application\_name***.
2. Under Related Items, click **Web Modules > *URI\_file\_name* > Web Services: Server Security Bindings**.
3. In the Request Receiver Binding column, click **Edit > Signing Information**.

To configure the trust anchor, complete the following steps:

1. Click **Servers > Application Servers > *server\_name***.
2. Under Additional Properties, click **Web Services: Default bindings for Web Services Security > Trust Anchors**.

To configure the collection certificate store, complete the following steps:

1. Click **Servers > Application Servers > *server\_name***.
2. Under Related Items, click **Web Services: Default bindings for Web Services Security > Collection Certificate Store**.

To configure the key locators, complete the following steps:

1. Click **Servers > Application Servers > *server\_name***.
2. Under Additional Properties, click **Web Services: Default bindings for Web Services Security > Key Locators**.

### Request decryption

If the confidentiality constraints (encryption) are specified, verify that the encryption information is defined.

To configure the encryption information parameters, complete the following steps:

1. Click **Enterprise Applications** > *application\_name*.
2. Under Related Items, click **Web Module**.
3. Under Additional Properties, click **Web Services: Server Security Bindings**. Under Request Receiver Binding, click **Edit** > **Encryption Information**.

To configure the key locators, complete the following steps:

1. Click **Servers** > **Application Servers** > *server\_name*.
2. Under Additional Properties, click **Web Services: Default bindings for Web Services Security** > **Key Locators**.

### BasicAuth authentication

If BasicAuth authentication is configured as the required security token, specify the CallbackHandler in the binding file to collect the basic authentication data. The following list contains Callback support implementations:

#### **com.ibm.wsspi.wssecurity.auth.callback.GuiPromptCallbackHandler**

The implementation prompts for BasicAuth information (user name and password) in an interface panel.

#### **com.ibm.wsspi.wssecurity.auth.callback.NonPromptCallbackHandler**

This implementation reads the BasicAuth information from the binding file.

#### **com.ibm.wsspi.wssecurity.auth.callback.StdPromptCallbackHandler**

This implementation prompts for a user name and password using the standard in (stdin) prompt.

To configure the login mapping information, complete the following steps:

1. Click **Server** > **Application Servers** > *server\_name*.
2. Under Additional Properties, click **Web Services: Default bindings for Web Services Security** > **Login Mappings**.

### Identity (ID) Assertion authentication with the BasicAuth TrustMode

Configure a login binding in the bindings file with a `com.ibm.wsspi.wssecurity.auth.callback.NonPromptCallbackHandler` implementation. Specify a BasicAuth user ID and password that a TrustedIDEvaluator on a downstream server trusts.

To configure the login mapping information, complete the following steps:

1. Click **Server** > **Application Servers** > *server\_name*.
2. Under Additional Properties, click **Web Services: Default bindings for Web Services Security** > **Login Mappings**.

### Identity (ID) Assertion authentication with the Signature TrustMode

Configure the signing information in the bindings file with a signing key pointing to a key locator. The key locator contains the X.509 certificate that is trusted by the downstream server.

To configure the login mapping information, complete the following steps:

1. Click **Server** > **Application Servers** > *server\_name*.
2. Under Additional Properties, click **Web Services: Default bindings for Web Services Security** > **Login Mappings**.

The Java Authentication and Authorization Service (JAAS) uses `WSLogin` as the name of the login configuration. To configure JAAS, click **Security** > **JAAS Configuration** > **Application Logins**.

The value of the <TrustedIDEvaluatorRef> tag in the binding must match the value of the <TrustedIDEvaluator> name.

To configure the trusted ID evaluators, complete the following steps:

1. Click **Servers > Application Servers > *server\_name***.
2. Under Additional Services, click **Web Services: Default bindings for Web Services Security > Trusted ID Evaluators**.

### Response signing

If the integrity constraints (digital signature) are defined, verify that you have the signing information configured in the binding files.

To specify the signing information, complete the following steps:

1. Click **Applications > Enterprise Applications > *application\_name***.
2. Under Related Items, click **Web Modules > *URI\_file\_name* > Web Services: Server Security Bindings**.
3. In the Request Receiver Binding column, click **Edit > Signing Information**.

To configure the key locators, complete the following steps:

1. Click **Servers > Application Servers > *server\_name***.
2. Under Additional Properties, click **Web Services: Default bindings for Web Services Security > Key Locators**.

### Response encryption

If the confidentiality constraints (encryption) are specified, verify that the encryption information is defined.

To specify the encryption information, complete the following steps:

1. Click **Enterprise Applications > *application\_name***.
2. Under Related Items, click **Web Module**.
3. Under Additional Properties, click **Web Services: Server Security Bindings**.
4. Under Request Receiver Binding, click **Edit > Encryption Information**.

To configure the key locators, complete the following steps:

1. Click **Servers > Application Servers > *server\_name***.
2. Under Additional Properties, click **Web Services: Default bindings for Web Services Security > Key Locators**.

## Authentication method overview

The Web services security implementation for WebSphere Application Server supports the following authentication methods: BasicAuth, Lightweight Third Party Authentication (LTPA), digital signature, and identity assertion.

When the WebSphere Application Server is configured to use the BasicAuth authentication method, the sender attaches the LTPA token as a BinarySecurityToken from the current security context or from basic authentication data configuration in the binding file in the SOAP message header. The Web services security message receiver authenticates the sender by validating the user name and password against the configured user registry. With the LTPA method, the sender attaches the LTPA BinarySecurityToken it previously received in the SOAP message header. The receiver authenticates the sender by validating the LTPA token and the token expiration time. With the Digital Signature authentication method, the sender attaches a BinarySecurityToken from a X509 certificate to the Web services security message header



along with a digital signature of the message body, time stamp, security token, or any combination of the three. The receiver authenticates the sender by verifying the validity of the X.509 certificate and the digital signature using the public key from the verified certificate.

The identity assertion authentication method is different from the other three authentication methods. This method establishes the security credential of the sender based on the trust relationship. You can use the identity assertion authentication method, for example, when an intermediary server must invoke a service from a downstream server on behalf of the client, but does not have the client authentication information. The intermediary server might establish a trust relationship with the downstream server and then assert the client identity to the same downstream server.

Web Services Security supports the following trust modes:

- BasicAuth
- Digital signature
- Presumed trust

When you use the BasicAuth and digital signature trust modes, the intermediary server passes its own authentication information to the downstream server for authentication. The presumed trust mode establishes a trust relationship using some external mechanism. For example, the intermediary server might pass SOAP messages through a Secure Socket Layers (SSL) connection with the downstream server and transport layer client certificate authentication.

The Web services security implementation for WebSphere Application Server validates the trust relationship by following this procedure:

1. The downstream server validates the authentication information of the intermediary server.
2. The downstream server verifies whether the authenticated intermediary server is authorized for identity assertion. For example, the intermediary server must be in the trust list for the downstream server.

The client identity might be represented by a name string, a distinguished name (DN), or an X.509 certificate. The client identity is attached in the Web services security message in a UsernameToken with just a user name, DN, or in a BinarySecurityToken of a certificate. The following table summarizes the type of security token that is required for each authentication method.

*Table 6. Authentication methods and their security tokens*

Authentication method	Security token
BasicAuth	BasicAuth requires <wsse:UsernameToken> with <wsse:Username> and <wsse:Password>.
Signature	Signature requires <ds:Signature> and <wsse:BinarySecurityToken>.
IDAssertion	IDAssertion requires <wsse:UsernameToken> with <wsse:Username> or <wsse:BinarySecurityToken> with a X.509 certificate for client identity depending on <idType>. This method also requires other security tokens according to the <trustMode>: <ul style="list-style-type: none"> <li>• If the &lt;trustMode&gt; is BasicAuth, IDAssertion requires &lt;wsse:UsernameToken&gt; with &lt;wsse:Username&gt; and &lt;wsse:Password&gt;.</li> <li>• If the &lt;trustMode&gt; is Signature, IDAssertion requires &lt;wsse:BinarySecurityToken&gt;.</li> </ul>
LTPA	LTPA requires <wsse:BinarySecurityToken> with an LTPA token.

A Web service can support multiple authentication methods simultaneously. The receiver side of the Web services deployment descriptor can specify all the authentication methods that are supported in the `ibm-webservices-ext.xml` XML file. The Web services receiver-side, as shown in the following example, is configured to accept all the authentication methods described previously:

```
<loginConfig xmi:id="LoginConfig_1052760331326">
 <authMethods xmi:id="AuthMethod_1052760331326" text="BasicAuth"/>
 <authMethods xmi:id="AuthMethod_1052760331327" text="IDAssertion"/>
 <authMethods xmi:id="AuthMethod_1052760331336" text="Signature"/>
 <authMethods xmi:id="AuthMethod_1052760331337" text="LTPA"/>
</loginConfig>
<idAssertion xmi:id="IDAssertion_1052760331336" idType="Username" trustMode="Signature"/>
```

You can define only one authentication method in the sender-side Web services deployment descriptor. A Web service client can use any of the authentication methods that are supported by the particular Web services application. The following example illustrates an identity assertion authentication method configuration in the `ibm-webservicesclient-ext.xml` deployment descriptor extension of the Web service client:

```
<loginConfig xmi:id="LoginConfig_1051555852697">
 <authMethods xmi:id="AuthMethod_1051555852698" text="IDAssertion"/>
</loginConfig>
<idAssertion xmi:id="IDAssertion_1051555852697" idType="Username" trustMode="Signature"/>
```

As shown in the previous example, the client identity type is Username and the trust mode is digital signature.

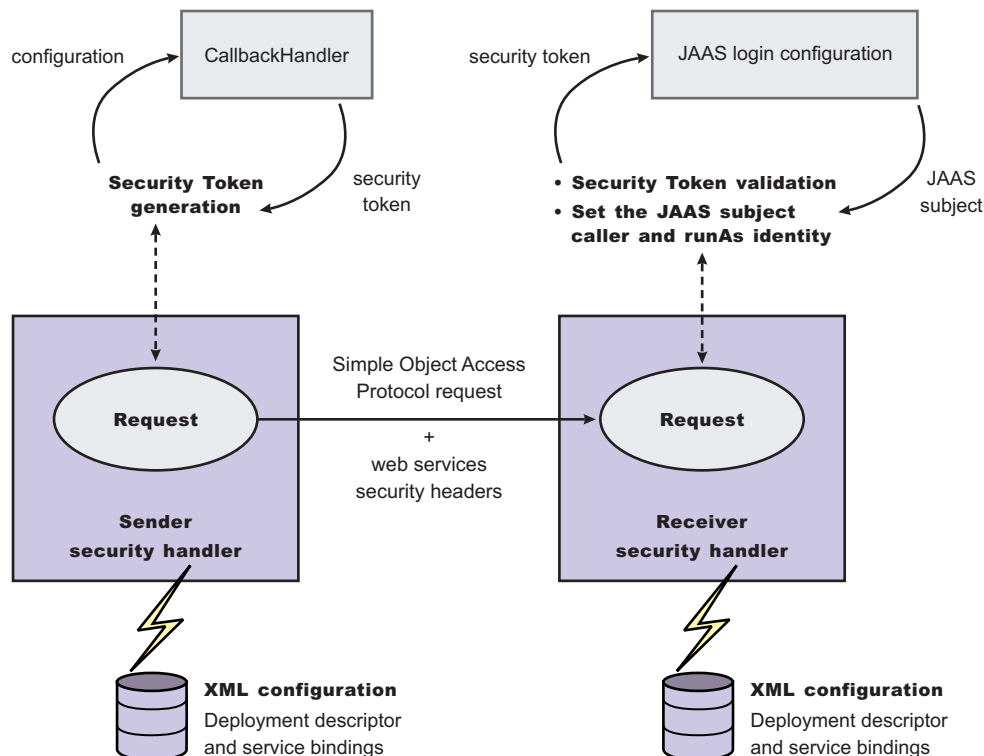


Figure 20. Security token generation and validation

The sender security handler invokes the `handle()` method of an implementation of the `javax.security.auth.callback.CallbackHandler` interface. The `javax.security.auth.callback.CallbackHandler` interface creates the security token and passes it back to

the sender security handler. The sender security handler constructs the security token based on the authentication information in the callback array and inserts the security token into the Web services security message header.

The receiver security handler compares the token type in the message header with the expected token types configured in the deployment descriptor. If none of the expected token types are found in the Web services security header of the SOAP message, the request is rejected with a SOAP fault exception. Otherwise, the token type is used to map to a Java Authentication and Authorization Service (JAAS) login configuration for validating the token. If the authentication is successful, a JAAS Subject is created and associated with the running thread. Otherwise, the request is rejected with a SOAP fault exception.

## XML digital signature

XML-Signature Syntax and Processing (XML signature) is a specification that defines XML syntax and processing rules to sign and verify digital signatures for digital content. The specification was developed jointly by the World Wide Web Consortium (W3C) and the Internet Engineering Task Force (IETF).

XML signature does not introduce new cryptographic algorithms. WebSphere Application Server uses XML signature with existing algorithms such as RSA, HMAC, and SHA1. XML signature defines many methods for describing key information and enables the definition of a new method.

XML canonicalization (c14n) is often needed when you use XML signature. Information can be represented in various ways within serialized XML documents. For example, although their octet representations are different, the following examples are identical:

- `<person first="John" last="Smith"/>`
- `<person last="Smith" first="John"></person>`

C14n is a process used to canonicalize XML information. Select an appropriate c14n algorithm because the information that is canonicalized is dependent upon this algorithm. One of the major c14n algorithms, Exclusive XML Canonicalization, canonicalizes the character encoding scheme, attribute order, namespace declarations, and so on. The algorithm does not canonicalize white space outside tags, namespace prefixes, or data type representation.

### XML signature in the Web Services Security-Core specification

The Web Services Security-Core (WSS-Core) specification defines a standard way for Simple Object Access Protocol (SOAP) messages to incorporate an XML signature. You can use almost all of the XML signature features in WSS-Core except enveloped signature and enveloping signature. However, WSS-Core has some recommendations such as exclusive canonicalization for the c14n algorithm and some additional features such as `SecurityTokenReference` and `KeyIdentifier`. The `KeyIdentifier` is the value of the `SubjectKeyIdentifier` field within the X.509 certificate. For more information on the `KeyIdentifier`, see "Reference to a Subject Key Identifier" within the OASIS Web Services Security X.509 Certificate Token Profile documentation.

By including XML signature in SOAP messages, the following are realized:

#### Message integrity

A message receiver can confirm that attackers or accidents have not altered parts of the message after these parts are signed by a key.

#### Authentication

You can assume that a valid signature is *proof of possession*. A message with a digital certificate issued by a certificate authority and a signature in the message that is validated successfully by a public key in the certificate, is proof that the signer has the corresponding private key. The receiver can authenticate the signer by checking the trustworthiness of the certificate.

## XML signature in the current implementation

XML signature is supported in Web services security, however, an application programming interface (API) is not available. The current implementation has many hardcoded behaviors and has some user-operable configuration items. To configure the client for digital signature, see [Configuring the client for response digital signature verification: Verifying the message parts](#). To configure the server for digital signature, see [Configuring the server for request digital signature verification: Verifying the message parts](#).

## Security considerations

In a replay attack, an attacker taps the lines, receives a signed message, and then returns the message to the receiver. In this case, the receiver receives the same message twice and might process both of them if the signatures are valid. Processing both messages can cause damage to the receiver if the message is a claim for money. If you have the signed generation time stamp and the signed expiration time in a message replay, attacks might be reduced. However, this is not a complete solution. A message must have a nonce value to prevent these attacks and the receiver must reject a message that contains a processed nonce. The current implementation does not provide a standard way to generate and check nonces in messages. In WebSphere Application Server, Version 5.1, nonce is supported in username tokens only. The username token profile contains concrete nonce usage scenarios for username tokens. Applications handle nonces (such as serial numbers) and they need to be signed.

## Signing information collection

Use this page to view a list of signing parameters. Signing information is used to sign and validate parts of a message including the body, time stamp, and user name token. You can also use these parameters for X.509 validation when the authentication method is IDAssertion and the ID type is X509Certificate in the server-level configuration. In such cases, you must fill in the certificate path fields only.

To view this administrative console page, complete the following steps:

1. Click **Applications** > **Enterprise Applications** > *application\_name*.
2. Under Related Items, click **Web Modules** > *URI\_file\_name* > **Web Services: Server Security Bindings**.
3. In the Request Receiver Binding column, click **Edit** > **Signing Information**.
4. Click **New** to create a signing parameter. Click **Delete** to delete a signing parameter.

### **Signature Method:**

Specifies the unique name of the signature method.

## Signing information configuration settings

Use this page to configure new signing parameters.

The specifications listed on this page for the signature method, digest method, and canonicalization method are located in the W3C document entitled, "XMLSignature Syntax and Specification: W3C Recommendation 12 Feb 2002".

To view this administrative console page:

1. Click **Enterprise Applications** > *application\_name*.
2. Under Related Items, click **Web Modules** > *URI\_file\_name* > **Web Services: Server Security Bindings**.
3. In the Request Receiver Binding column, click **Edit** > **Signing Information**.
4. Click **New** to create a signing parameter or click **Delete** to delete a signing parameter.

### **Signature Method:**

Specifies the algorithm Uniform Resource Identifiers (URI) of the signature method. This method contains the actual value of the digital signature encoded using base64.

The following algorithms are supported:

- <http://www.w3.org/2000/09/xmlsig#rsa-sha1>
- <http://www.w3.org/2000/09/xmlsig#dsa-sha1>

#### **Digest Method:**

Specifies the algorithm URI of the digest method.

The <http://www.w3.org/2000/09/xmlsig#sha1> algorithm is supported.

#### **Canonicalization Method:**

Specifies the algorithm URI of the canonicalization method.

The following algorithms are supported:

- <http://www.w3.org/2001/10/xml-exc-c14n#>
- <http://www.w3.org/2001/10/xml-exc-c14n#WithComments>
- <http://www.w3.org/TR/2001/REC-xml-c14n-20010315>
- <http://www.w3.org/TR/2001/REC-xml-c14n-20010315#WithComments>

#### **Signing Key:**

Specifies the key information that is used for signing. These fields are ignored in receiver-side configuration.

If you specify a **Key Name** and a **Key Locator Reference**, select **None** for the Certificate Path.

#### **Certificate Path:**

Specifies the settings for the certificate path validation. When you select **Trust Any**, this validation is skipped and all the incoming certificates are trusted. These fields are ignored in sender-side configuration.

If you click **Trust Any** or select a **Trust Anchor** and a **Certificate Store**, select **None** for the Signing Key in the previous field.

#### **Trust Anchor**

The selections available for Trust Anchor are specified by clicking **Servers > Application Servers > *server\_name***. Under Additional Properties, click **Web Services: Default bindings for Web Services Security > Trust Anchors**.

#### **Certificate Store**

The selections available for the Collection Store are specified by clicking **Servers > Application Servers > *server\_name***. Under Related Items, click **Web Services: Default bindings for Web Services Security > Collection Certificate Store**.

### **Signing parameter configuration settings**

Use this page to configure new signing parameters.

The specifications listed on this page for the signature method, digest method, and canonicalization method are located in the World Wide Web Consortium (W3C) document entitled, *XML Signature Syntax and Specification: W3C Recommendation 12 Feb 2002*.

To view this administrative console page, complete the following steps:

1. Click **Enterprise Applications** > *application\_name*.
2. Under Related Items, click **Web Modules** > *URI\_file\_name* > **Web Services: Client Security Bindings**.
3. In the Request Sender Binding column, click **Edit** > **Signing Information**.

If the signing information is not available, select **None**.

If the signing information is available, select **Dedicated Signing Information** and specify the configuration in the following fields:

***Signature Method:***

Specifies the algorithm Uniform Resource Identifiers (URI) of the signature method. This method contains the actual value of the digital signature encoded using base64.

The following algorithms are supported:

- <http://www.w3.org/2000/09/xmlsig#rsa-sha1>
- <http://www.w3.org/2000/09/xmlsig#dsa-sha1>

***Digest Method:***

Specifies the algorithm URI of the digest method.

The <http://www.w3.org/2000/09/xmlsig#sha1> algorithm is supported.

***Canonicalization Method:***

Specifies the algorithm URI of the canonicalization method.

The following algorithms are supported:

- <http://www.w3.org/2001/10/xml-exc-c14n#>
- <http://www.w3.org/2001/10/xml-exc-c14n#WithComments>
- <http://www.w3.org/TR/2001/REC-xml-c14n-20010315>
- <http://www.w3.org/TR/2001/REC-xml-c14n-20010315#WithComments>

***Signing Key:***

Specifies the key information that is used for signing. These fields are ignored in receiver-side configuration.

If the signing key is not available, select **None**.

***Certificate Path:***

Specifies the settings for the certificate path validation. When you select **Trust Any**, this validation is skipped and all the incoming certificates are trusted. These fields are ignored in sender-side configuration.

If there is not a certificate path, select **None**.

If there is a certificate path, select **Trust Any** or select a Trust Anchor and a Certificate Store.

**Trust Anchor**

Specify the selections for the Trust Anchor field by clicking **Servers > Application Servers > *server\_name***. Under Additional Properties, click **Web Services: Default bindings for Web Services Security > Trust Anchors**.

### Certificate Store

Specify the selections for the Collection Store field by clicking **Servers > Application Servers *server\_name***. Under Related Items, click **Web Services: Default bindings for Web Services Security > Collection Certificate Store**.

## Securing Web services using XML digital signature

WebSphere Application Server provides several different methods to secure your Web services; Extensible Markup Language (XML) digital signature is one of these methods. You might secure your Web services using any of the following methods:

- XML digital signature
- XML encryption
- Basicauth authentication
- Identity assertion authentication
- Signature authentication
- Pluggable token

XML digital signature provides both message integrity and authentication capabilities when it is used with SOAP messages. A message receiver can verify that attackers or accidents have not altered parts of the message after the message was signed by a key. If a message has a digital certificate issued by a certificate authority (CA) and a signature in the message is validated successfully by a public key in the certificate, it is proof that the signer has the corresponding private key. To use XML digital signature to secure Web services, complete the following steps:

1. Define the security constraints or extensions. To configure the security constraints, you must use the Application Server Toolkit, which is available at the following Web site:  
[http://www.ibm.com/support/docview.wss?rs=180&context=SSEQTP&q=ASTK&uid=swg24005125&loc=en\\_US&cs=utf-8&lang=en+en](http://www.ibm.com/support/docview.wss?rs=180&context=SSEQTP&q=ASTK&uid=swg24005125&loc=en_US&cs=utf-8&lang=en+en)
  - a. Configure the client to digitally sign a message request. To configure the client, complete the following steps to specify which parts of the SOAP message to digitally sign and define the method used to digitally sign the message. The client in these steps is the request sender.
    - 1) Specify the message parts by following the steps found in Configuring the client for request signing: digitally signing message parts.
    - 2) Select the method used to digitally sign the request message. You can select the digital signature method by following the steps in Configuring the client for request signing: choosing the digital signature method.
  - b. Configure the server to verify the digital signature that is used in the message request. To configure the server, you must specify which parts of the SOAP message, sent by the request sender, contain digitally signed information and which method was used to digitally sign the message. The settings chosen for the request receiver, or the server in this step, must match the settings chosen for the request sender in the previous step.
    - 1) Define the message parts by following the steps found in Configuring the server for request digital signature verification: verifying message parts.
    - 2) Select the same method used by the request sender to digitally sign the message. You can select the digital signature method by following the steps in Configuring the server for request digital signature verification: choosing the verification method



- c. Configure the server to digitally sign a message response. To configure the server, complete the following steps to specify which parts of the SOAP message to digitally sign and define the method used to digitally sign the message. The sender in these steps is the response sender.
    - 1) Specify which message parts to digitally sign by following the steps found in Configuring the server for response signing: digitally signing message parts.
    - 2) Select the method used to digitally sign the response message. You can select the digital signature method by following the steps in Configuring the server for response signing: choosing the digital signature method
  - d. Configure the client to verify the digital signature that is used in the message response. To configure the client, you must specify which parts of the SOAP message sent by the response sender contain digitally signed information and which method was used to digitally sign the message. The settings chosen for the response receiver, or client in this step, must match the settings chosen for the response sender in the previous step.
    - 1) Define the message parts by following the steps found in Configuring the client for response digital signature verification: verifying message parts
    - 2) Select the same method used by the response sender to digitally sign the message. You can select the digital signature method by following the steps in Configuring the client for response digital signature verification: choosing the verification method
2. Define the client security bindings. To configure the client security bindings, complete the steps in either of the following topics:
    - Configuring the client security bindings using the Application Server Toolkit
    - Configuring the client security bindings using the administrative console
  3. Define the server security bindings. To configure the server security bindings, complete the steps in either of the following topics:
    - Configuring the server security bindings using the Application Server Toolkit
    - Configuring the server security bindings using the administrative console

After completing these steps, you have secured your Web services using XML digital signature.

## Transport level security

*Transport level security* is based on Secure Sockets Layer (SSL) or Transport Layer Security (TLS) that runs beneath HTTP.

SSL and TLS provide security features including authentication, data protection, and cryptographic token support for secure HTTP connections. To run with HTTPS, the service endpoint address must be in the form `https://`.

The integrity and confidentiality of transport data, including SOAP messages and HTTP basic authentication, is confirmed when you use SSL and TLS. See "Secure Sockets Layer" in the Information Center. Web services applications can also use Federal Information Processing Standard (FIPS) approved ciphers for more secure TLS connections.

WebSphere Application Server uses the Java Secure Sockets Extension (JSSE) package to support SSL and TLS.

## Configuring HTTP outbound transport level security

The HTTP outbound transport-level security for a Web service is based on the Secure Sockets Layer (SSL) configuration of the WebSphere Application Server Web container. Review "Configuring Secure Sockets Layer" and "HTTP transport collection" in the Information Center.

To configure HTTP outbound transport-level security, complete the following steps:

1. Configure the HTTP outbound transport-level security settings.

- Use either the WebSphere Application Server administrative console or the Assembly Toolkit to configure the HTTP outbound transport-level security for a Web Service acting as a client to another Web service.
  - Before installing the Web services application, use the Assembly Toolkit to configure the HTTP SSL Configuration in the Web Services Client Port Binding page. The Web Services Client Port Binding page is available after double-clicking the `webservicesclient.xml` file.
  - After installing the Web services application, use the administrative console to configure the Web services client security binding collection. To access the collection, complete the following steps:
    - a. Click **Applications > Enterprise Applications**.
    - b. Under Related Items, click either **Web Modules** or **EJB Modules**.
    - c. Click the name of the URI.
    - d. Under Additional Properties, click **Web Services: Client Security Bindings**.

**Attention:** If the HTTP outbound transport-level security settings are not configured, the default Secure Sockets Layer (SSL) settings for the Java Secure Socket Extension (JSSE) file are used.

- Use the properties to configure the HTTP outbound transport-level security for a Web service client.
  - a. Create a property file that includes the following properties:

```
com.ibm.ssl.protocol
com.ibm.ssl.keyStoreType
com.ibm.ssl.keyStore
com.ibm.ssl.keyStorePassword
com.ibm.ssl.trustStoreType
com.ibm.ssl.trustStore
com.ibm.ssl.trustStorePassword
```

- b. Set the `com.ibm.webservices.sslConfigURL` Java system property to the absolute path of the created property file.

**Attention:** If the outbound transport-level security is not configured, the default SSL settings of the JSSE file are used.

2. Optional: Accept the redirection of HTTP request to a different URI in HTTPS. A redirection of the HTTP request to a different URI in HTTPS can occur if the transport guarantee of CONFIDENTIAL or INTEGRAL is configured in the application. To accept the redirection, you can do either of the following tasks:
  - Set the `com.ibm.ws.webservices.HttpRedirectEnabled` Java system property to `true`.
  - Programmatically set the `com.ibm.wsspi.webservices.Constants.HTTP_REDIRECT_ENABLED` property to `true` in the stub or call object before invoking the service.

### **HTTP SSL Configuration collection:**

Use this page to configure transport-level Secure Sockets Layer (SSL) security. You can use this configuration when a Web service is a client to another Web service.

You can use transport-level security to enable HTTP SSL (or HTTPS). Transport-level security can be enabled or disabled independently from message-level security. Because transport-level security provides minimal security, use message-level security when security is essential to the Web service application.

To view this administrative console page, complete the following steps:

1. Click **Applications > Enterprise Applications > *application\_name***.
2. Under Related Items, click **Web Module > *URI\_file\_name* > Web Services: Client Security Bindings**.
3. Under HTTP SSL Configuration, click **Edit**.

### *HTTP SSL Enabled:*

Specifies secure socket communications for the HTTP transport for this port. When enabled, WebSphere Application Server uses the HTTP SSL Configuration setting.

### *HTTP SSL Configuration:*

Specifies which alias of the SSL configuration to use with the HTTP transport for this port.

This option is used if you select **HTTP SSL Enabled**. SSL aliases are defined in the Secure Sockets Layer configuration repertoire, which you can configure by clicking **Security > SSL**.

## **HTTP basic authentication**

*HTTP basic authentication* uses a user name and password to authenticate a service client to a secure endpoint.

WebSphere Application Server can have several resources, including Web services, protected by a Java 2 Platform, Enterprise Edition (J2EE) security model.

A simple way to provide authentication data for the service client is to authenticate to the protected service endpoint to the HTTP basic authentication. The basic authentication is located in the HTTP header that carries the SOAP request. When the application server receives the HTTP request, the user name and password are retrieved and verified using the authentication mechanism specific to the server.

Although the basic authentication data is base64-encoded, sending data over HTTPS is recommended. The integrity and confidentiality of the data can be protected by the Secure Sockets Layer (SSL) protocol.

In some cases, a firewall is present using the pass-thru HTTP proxy server. The HTTP proxy server forwards the basic authentication data into the J2EE application server. The proxy server can also be protected. Applications can specify the proxy data by setting properties in a stub object.

## **Configuring HTTP basic authentication**

To configure HTTP basic authentication, use the WebSphere Application Server tools to modify the binding information or programmatically set properties in a Stub or Call object. The values set programmatically take precedence over the values defined in the binding. However, you only can programmatically configure the HTTP proxy authentication.

1. Use either the Assembly Toolkit or the administrative console to configure the HTTP basic authentication
  - Before installing the Web services application, use Assembly toolkit to configure the HTTP basic authentication in the Web Services Client Port Binding page for a Web service or a Web service client. The Web Services Client Port Binding page is available after double-clicking the `webservicesclient.xml` file.
  - After installing the Web services application, you can use administrative console to configure the Web services client security bindings for a Web service only. To access the collection, complete the following steps:
    - a. Click **Applications > Enterprise Applications**.
    - b. Under Related Items, click either **Web Modules** or **EJB Modules**.
    - c. Click the name of the URI.
    - d. Under Additional Properties, click **Web Services: Client Security Bindings**.
2. Programmatically set the properties in the stub or call object for a Web service or a Web service client  
Programmatically set the following properties:

```
javax.xml.rpc.Call.USERNAME_PROPERTY
javax.xml.rpc.Call.PASSWORD_PROPERTY
javax.xml.rpc.Stub.USERNAME_PROPERTY
javax.xml.rpc.Stub.PASSWORD_PROPERTY
```

3. Programmatically set the properties in the stub or call object to configure the HTTP proxy authentication. Programmatically set the following properties for HTTP:

```
com.ibm.wsspi.webservices.HTTP_PROXYHOST_PROPERTY
com.ibm.wsspi.webservices.HTTP_PROXYPORT_PROPERTY
com.ibm.wsspi.webservices.HTTP_PROXYUSER_PROPERTY
com.ibm.wsspi.webservices.HTTP_PROXYPASSWORD_PROPERTY
```

Programmatically set the following properties for HTTPS:

```
com.ibm.wsspi.webservices.HTTPS_PROXYHOST_PROPERTY
com.ibm.wsspi.webservices.HTTPS_PROXYPORT_PROPERTY
com.ibm.wsspi.webservices.HTTPS_PROXYUSER_PROPERTY
com.ibm.wsspi.webservices.HTTPS_PROXYPASSWORD_PROPERTY
```

### ***HTTP basic authentication collection:***

Use this page to specify a user ID and password for transport-level basic authentication security for this port. You can use this configuration when a Web service is a client to another Web service.

You can use transport-level security to enable basic authentication. Transport-level security can be enabled or disabled independently from message-level security. Because transport-level security provides minimal security, use message-level security when security is essential to the Web service application.

To view this administrative console page, complete the following steps:

1. Click **Applications > Enterprise Applications > application\_name**.
2. Under Related Items, click **Web Module > URI\_file\_name > Web Services: Client Security Bindings**.
3. Under HTTP Basic Authentication, click **Edit**.

#### *Basic Authentication ID:*

Specifies the user ID for the HTTP basic authentication for this port.

#### *Basic Authentication Password:*

Specifies the password for the HTTP basic authentication for this port.

## **Default configuration for WebSphere Application Server**

In the WebSphere Application Server, each application server has a copy of the `ws-security.xml` file, which defines the default binding information for Web services security. The following list contains the defaults defined in the `ws-security.xml` file:

### **Trust anchors**

Identifies the trusted root certificates for signature verification.

### **Collection certificate stores**

Contains certificate revocation lists (CRLs) and nontrusted certificates for verification.

### **Key locators**

Locates the keys for digital signature and encryption.

### Trusted ID evaluators

Evaluates the trust of the received identity before identity assertion.

### Login mappings

Contains the Java Authentication and Authorization Service (JAAS) configurations for AuthMethod token validation.

If the Web services security constraints specified in the deployment descriptors and the required bindings are not defined in the bindings file, the default constraints in the `ws-security.xml` file are used.

When you use the **addNode** command, the `ws-security.xml` file is added with the server configuration to the new cell. The following figure shows the activity when you use the **addNode** command.

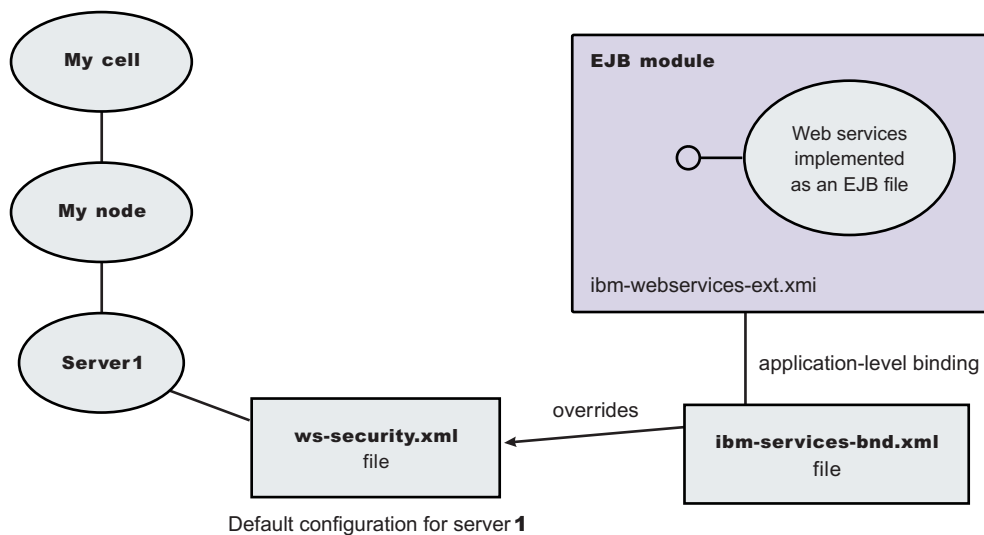


Figure 21. Configuration when using the **addNode** command

### Default binding:

The default binding information is defined in the `ws-security.xml` file and can be administered by either the administrative console or by scripting.

Certain applications can share certain binding information. This information includes truststores, keystores, and authentication methods (token validation). WebSphere Application Server provides support for default binding information. Administrators can define binding information at the server level and applications can refer to this binding information.

You can define the following binding information in the `ws-security.xml` file:

### Trust anchors (truststore)

- *Trust anchors* contain key store configuration information that has the root-trusted certificates. Trust anchors are used for certificate path validation of the incoming X.509-formatted security tokens.
- The Trust Anchor Name is used in the binding file (`ibm-webservices-bnd.xmi` and `ibm-webservicesclient-bnd-xmi` when Web services is running as a client) to refer to the trust anchor defined in the default binding information. The trust anchor name must be unique in the trust anchor collection.

## Collection certificate store

- The *collection certificate store* specifies a list of untrusted, intermediate certificates and is used for certificate path validation of incoming X.509-formatted security tokens. The default provider is IBM CertPath.
- The Certificate Store Name is used in the binding file (`ibm-webservices-bnd.xml` and `ibm-webservicesclient-bnd.xml` when Web services is running as a client) to refer to the certificate store defined in the default binding information. The Certificate Store Name must be unique to the collection certificate store collection.

## Key locators

- *Key locators* specify implementation of the `com.ibm.wsspi.wssecurity.config.KeyLocator` interface. This interface is used to retrieve keys for signature or encryption. Customer implementations can extend the key locator interface to retrieve keys using other methods. WebSphere Application Server provides implementations to retrieve a key from the key store, map an authenticated identity to a key in the key store, or retrieve a key from the signer certificate (mapping and retrieving actions are used for encrypting the response).
- The Key Locator Name is used in the binding file (`ibm-webservices-bnd.xml` and `ibm-webservicesclient-bnd.xml` when Web services is running as a client) to refer to the key locator defined in the default binding information. The Key Locator Name must be unique to the key locators collection in the default binding information.

## Trusted ID evaluators

- *Trusted ID evaluators* are an implementation of the `com.ibm.wsspi.wssecurity.id.TrustedIDEvaluator` interface. This interface is used to make sure the identity (ID)-asserting authority is trusted. Additionally, you can extend the trusted identity evaluator to validate the trust. WebSphere Application Server provides a default implementation for validating trust based on a predefined list of identities.
- The Trusted ID Evaluator Name is used in the binding file (`ibm-webservices-bnd.xml`) to refer to the trusted identity evaluator defined in the default binding information. The Trusted ID Evaluator Name must be unique to the Trusted ID Evaluator collection.

## Login mappings

- *Login mappings* define the mapping of the authentication method to the Java Authentication and Authorization Service (JAAS) login configuration. The mappings are used to authenticate the incoming security token embedded in the Web services security Simple Object Access Protocol (SOAP) message header. The JAAS login configuration is defined in the administrative console under **Security > JAAS Configuration > Application Logins**.
- WebSphere Application Server defines the following authentication methods:

### BasicAuth

Authenticates user name and password.

### Signature

Maps the subject distinguished name (DN) in the certificate to a WebSphere Application Server credential.

### IDAssertion

Maps the identity to a WebSphere Application Server credential.

**LTPA** Authenticates a Lightweight Third Party Authentication (LTPA) token.

After identity authentication, the associated credential is used in the downstream call.

- This method can be extended to authenticate custom security tokens by providing a custom JAAS login configuration and by using the `com.ibm.wsspi.wssecurity.auth.module.WSSecurityMappingModule` to create the principal and credential required by WebSphere Application Server.

- If LoginConfig (AuthMethod) is defined in the IBM extension deployment descriptor (`ibm-webservices-ext.xmi`), but there are no login mapping bindings (`ibm-webservices-bnd.xmi`) defined for the AuthMethod, Web services security run time uses the login mapping defined in the default binding information.

## WebSphere Application Server

In the WebSphere Application Server, each server has a copy of the `ws-security.xml` file (default binding information for Web services security). There is no cell-level copy of the `ws-security.xml` file, which is only available in the WebSphere Application Server Network Deployment installation. To navigate to the server-level default binding in the administrative console, complete the following steps:

1. Click **Servers > Application Servers > server1**.
2. Under Additional Properties, click **Web Services: Default bindings for Web Services Security**.

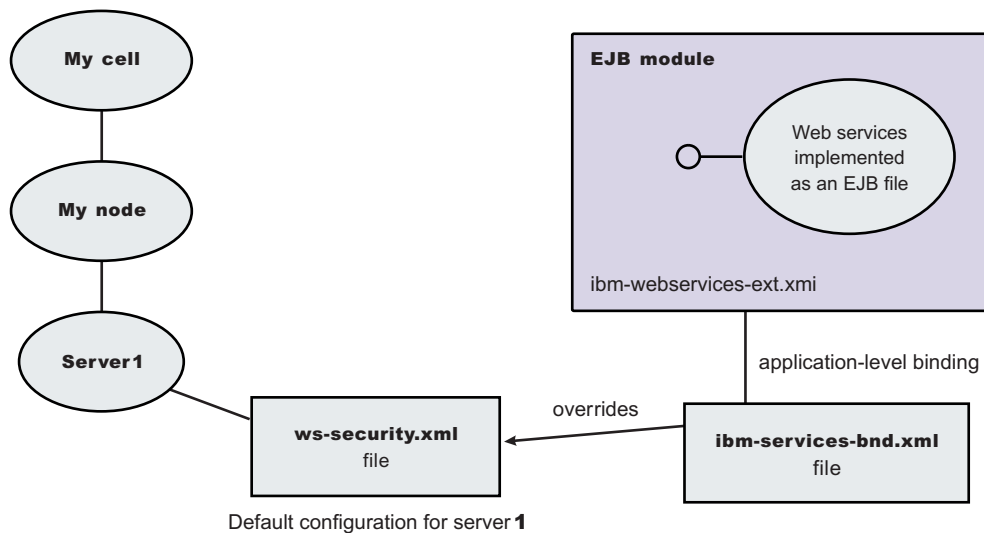


Figure 22. Web services security application-level bindings and server-level default binding information

Web services security run time uses the binding information in the application Enterprise JavaBeans (EJB) or Web module binding file (`ibm-webservices-bnd.xmi` or `ibm-webservicesclient-bnd.xmi` if Web services is acting as a client on the server) if the binding information is defined in the application-level binding file. For example, if key locator K1 is defined in both the application-level binding file and the default binding file (`ws-security.xml`), the K1 in the application-level binding file is used.

*Web services: default bindings for the Web services security collection:*

Use this page to configure the settings for nonce on the server level and to manage the default bindings for trust anchors, the collection certificate store, key locators, trusted ID evaluators, and login mappings.

To view this administrative console page, click **Servers > Application Servers > server\_name**. Under Additional Properties, click **Web Services: Default bindings for Web Services Security**.

Read the Web services documentation before you begin defining the default bindings for Web services security.

To define the server bindings, complete the following steps:



1. Click **Applications > Enterprise Applications > *application\_name***.
2. Under Related Items, click **Web Modules > *URI\_file\_name* > Web Services: Server Security Bindings**.

To define the client bindings, complete the following steps:

1. Click **Applications > Enterprise Applications > *application\_name***.
2. Under Related Items, click **Web Modules > *URI\_file\_name* > Web Services: Client Security Bindings**.

**5.1 +** *Nonce* is a unique cryptographic number embedded in a message to help stop repeat, unauthorized attacks of user name tokens. In a base WebSphere Application Server environment, you must specify values for the **Nonce Cache Timeout**, **Nonce Maximum Age**, and **Nonce Clock Skew** fields for the server-level.

The default binding configuration provides a central location where reusable binding information is defined. The application binding file can reference the information contained in the default binding configuration.

*Nonce Cache Timeout:*

Specifies the timeout value, in seconds, for the nonce cached on the server. Nonce is a randomly generated value.

The Nonce Cache Timeout field is required for the base WebSphere Application Server environment.

If you make changes to the nonce cache timeout value, you must restart WebSphere Application Server for the changes to take effect.

<b>Default</b>	600 seconds
<b>Minimum</b>	300 seconds

*Nonce Maximum Age:*

Specifies the default time, in seconds, before the nonce time stamp expires. Nonce is a randomly generated value.

The maximum value cannot exceed the number of seconds specified in the Nonce Cache Timeout field for the server level. The value set for this server-level Nonce Maximum Age field must not exceed Nonce Maximum Age value set for the cell level, which you can access by clicking **Security > Web Services > Properties**.

The Nonce Maximum Age field is required for the base WebSphere Application Server environment.

<b>Default</b>	300 seconds
<b>Range</b>	300 to Nonce Cache Timeout seconds

*Nonce Clock Skew:*

Specifies the default clock skew value, in seconds, to consider when WebSphere Application Server checks the timeliness of the message. Nonce is a randomly generated value.

The maximum value cannot exceed the number of seconds specified in the Nonce Maximum Age field.

The Nonce Clock Skew field is required for the base WebSphere Application Server environment.

<b>Default</b>	0 seconds
<b>Range</b>	0 to Nonce Maximum Age seconds

*Trust Anchors:*

Specifies a list of keystore objects that contain the trusted root certificates, self-signed or issued by a certificate authority (CA).

The certificate authority authenticates a user and issues a certificate. After the certificate is issued, the keystore objects, which contain these certificates, use the certificate for certificate path or certificate chain validation of incoming X.509-formatted security tokens.

*Collection Certificate Store:*

Specifies a list of the untrusted, intermediate certificate files.

The collection certificate store contains a chain of untrusted, intermediate certificates. The CertPath API attempts to validate these certificates, which are based on the trust anchor.

*Key Locators:*

Specifies a list of key locator objects that retrieves the keys for digital signature and encryption from a keystore file or a repository. The key locator maps a name or logical name to an alias or maps an authenticated identity to a key. This logical name is used to locate a key in a key locator implementation.

*Trusted ID Evaluators:*

Specifies a list of trusted ID evaluators that determines whether to trust the identity-asserting authority or the message sender.

The trusted ID evaluators are used to authenticate additional identities from one server to another server. For example, a client sends the identity of user A to server 1 for authentication. Server 1 calls downstream to server 2, asserts the identity of user A, and includes the user ID and password of server 1. Server 2 attempts to establish trust with server 1 by authenticating its user ID and password and checking the trust based on the TrustedIDEvaluator implementation. If the authentication process and the trust check are successful, server 2 trusts that server 1 authenticated user A and a credential is created for user A on server 2 to invoke the request.

*Login Mappings:*

Specifies a list of configurations for validating tokens within incoming messages.

Login mappings map the authentication method to the Java Authentication and Authorization Service (JAAS) configuration.

To configure JAAS, use the administrative console and click **Security > JAAS Configuration**.

*Web Services: Client security bindings collection:*

Use this page to view a list of client-side binding configurations for Web services security. These bindings are used when a Web service is a client to another Web service.

To view this administrative console page, complete the following steps:

1. Click **Enterprise Applications > application\_name**.

2. Under Related Items, click **Web Module** > *URI\_file\_name* > **Web Services: Client Security Bindings**.

*Port:*

Specifies the port used to send and receive messages from a server.

*Web Service:*

Specifies the name of the Web service.

*Request Sender Binding:*

Specifies the binding configuration used to send request messages to the request receiver.

Click **Edit** to configure the signing information, encryption information, and the login bindings for the request sender and to view a listing of key locators in the key store.

The binding information for the request sender that is specified for the client must match the binding information for the request receiver that is specified for the server.

*Response Receiver Binding:*

Specifies the binding configuration used to receive response messages from the response sender.

Click **Edit** to configure the signing and encryption information, and to view a list of trust anchors, intermediate certificates found in the collection certificate store, and the key locator objects for the response receiver.

The binding information for the response receiver that is specified for the client must match the binding information for the response sender that is specified for the server.

*HTTP Basic Authentication:*

Specifies the user ID and password to use for this port with HTTP transport-level basic authentication. You can enable transport-level authentication security independently of message-level security.

Click **Edit** to configure the basic authentication ID and password for transport-level authentication.

*HTTP SSL Configuration:*

Enables and configures transport-level Secure Socket Layer (SSL) security for this port. You can enable transport-level SSL security independently of message-level security.

Click **Edit** to specify the settings for transport-level HTTP SSL configuration for this port.

*Web services: Server security bindings collection:*

Use this page to view a list of server-side binding configurations for Web services security.

To view this administrative console page, complete the following steps:

1. Click **Enterprise Applications** > *application\_name*.
2. Under Related Items, click **Web Module** > *URI\_file\_name* > **Web Services: Server Security Bindings**.

*Port:*

Specifies the port in which messages are received from the request sender.

*Web Service:*

Specifies the name of the Web service.

*Request Receiver Binding:*

Specifies the binding configuration used to receive request messages from the request sender.

Click **Edit** to configure the signing information and encryption information and view a listing of trust anchors, intermediate certificates in the collection certificate store, key locators, trusted ID evaluators, and login mappings.

The binding information (for the request receiver) specified for the server must match the binding information (for the request sender) specified for the client.

*Response Sender Binding:*

Specifies the binding configuration used to send request messages to the response receiver.

Click **Edit** to configure the signing and encryption information, and to view a list of key locator objects for the response sender.

The binding information (for the response sender) specified for the server must match the binding information (for the response receiver) specified for the client.

## **Trust anchors**

A *trust anchor* specifies key stores that contain trusted root certificates that validate the signer certificate.

These key stores are used by the request receiver (as defined in the `ibm-webservices-bnd.xmi` file) and the response receiver (as defined in the `ibm-webservicesclient-bnd.xmi` file when Web services is acting as client) to validate the signer certificate of the digital signature. The key stores are critical to the integrity of the digital signature validation. If they are tampered with, the result of the digital signature verification is doubtful and comprised. Therefore, it is recommended that you secure these key stores. The binding configuration specified for the request receiver in the `ibm-webservices-bnd.xmi` file must match the binding configuration for the response receiver in the `ibm-webservicesclient-bnd.xmi` file.

The trust anchor is defined as `javax.security.cert.TrustAnchor` in the Java CertPath application programming interface (API). The Java CertPath API uses the trust anchor and the certificate store to validate the incoming X.509 certificate that is embedded in the SOAP message.

The Web services security implementation in WebSphere Application Server supports this trust anchor. In WebSphere Application Server, the trust anchor is represented as a Java key store object. The type, path, and password of the key store are passed to the implementation through the administrative console or by scripting.

## **Configuring trust anchors using the Assembly Toolkit**

This document describes how to configure trust anchors or trust stores at the application level. It does not describe how to configure trust anchors at the server or cell level. Trust anchors defined at the application level have a higher precedence over trust anchors defined at the server or cell level. You can configure an application-level trust anchor using the Assembly Toolkit or the administrative console. This document describes how to configure the application-level trust anchor using the Assembly Toolkit. For more information on creating and configuring trust anchors at the server or cell level, see either *Configuring the server security bindings using the Assembly Toolkit* or *Configuring the server security bindings using the administrative console*.

A trust anchor specifies key stores that contain trusted root certificates, which validate the signer certificate. These key stores are used by the request receiver (as defined in the `ibm-webservices-bnd.xml` file) and the response receiver (as defined in the `ibm-webservicesclient-bnd.xml` file when Web services is acting as client) to validate the signer certificate of the digital signature. The key stores are critical to the integrity of the digital signature validation. If they are tampered with, the result of the digital signature verification is doubtful and comprised. Therefore, it is recommended that you secure these key stores. The binding configuration specified for the request receiver in the `ibm-webservices-bnd.xml` file must match the binding configuration for the response receiver in the `ibm-webservicesclient-bnd.xml` file.

The steps in this document assume that you have already created a Web services-enabled Java 2 Platform, Enterprise Edition (J2EE) with Java Specification Requests (JSR) 109 enterprise application. If you have not created a Web services-enabled J2EE with JSR 109 enterprise application, see *Developing Web services*. Also, see either *Configuring the server security bindings using the Assembly Toolkit* or *Configuring the server security bindings using the administrative console* for an introduction on how to manage Web services security binding information on the server.

1. Configure the client-side response receiver, which is defined in the `ibm-webservicesclient-bnd.xml` bindings extensions file.
  - a. Launch the Assembly Toolkit and click **Windows > Open Perspective > J2EE**.
  - b. Select the Web services-enabled Enterprise JavaBeans (EJB) or Web module.
  - c. In the Package Explorer window, click the META-INF directory for an EJB module or the WEB-INF directory for a Web module.
  - d. Right-click the `webservicesclient.xml` file, select **Open With > Web Services Client Editor**, and click the **Web Services Client Binding** tab. The Web Services Client Binding editor is displayed.
  - e. Locate the Port Qualified Name Binding section and either select an existing entry or click **Add**, to add a new port binding. The Web Services Client Port Binding editor displays for the selected port.
  - f. Locate the Trust Anchor section and click **Add**. The Trust Anchor dialog box is displayed.
    - 1) Enter a unique name within the port binding for the **Trust anchor name**.  
The name is used to reference the trust anchor that is defined.
    - 2) Enter the key store password, path, and key store type.  
The supported key store types are Java Cryptography Extension (JCE) and JCEKS.Click **Edit** to edit the selected trust anchor.  
Click **Remove** to remove the selected trust anchor.  
When you start the application, the configuration is validated in the run time while the binding information is loading.
  - g. Save the changes.
2. Configure the server-side request receiver, which is defined in the `ibm-webservices-bnd.xml` bindings extensions file.
  - a. Launch the Assembly Toolkit and click **Windows > Open Perspective > J2EE**.
  - b. Select the Web services enabled EJB or Web module.
  - c. In the Package Explorer window, click the META-INF directory for an EJB module or the WEB-INF directory for a Web module.
  - d. Right-click the `webservices.xml` file, select **Open With > Web Services Editor**, and click the **Bindings** tab. The Web Services Bindings editor is displayed.
  - e. Locate the Web Service Description Bindings section and either select an existing entry or click **Add** to add a new Web services descriptor.
  - f. Click **Binding Configurations**. The Web Services Binding Configurations editor is displayed for the selected Web services descriptor.
  - g. Locate the Trust Anchor section and click **Add**. The Trust Anchor dialog box is displayed.
    - 1) Enter a unique name within the binding for the **Trust anchor name**.  
This unique name is used to reference the trust anchor defined.

- 2) Enter the key store password, path, and key store type. The supported key store types are JCE and JCEKS.

Click **Edit** to edit the selected trust anchor.

Click **Remove** to remove the selected trust anchor.

When you start the application, the configuration is validated in the run time while the binding information is loading.

- h. Save the changes.

This procedure defines trust anchors that can be used by the request receiver or the response receiver (if the Web services is acting as client) to verify the signer certificate.

The request receiver or the response receiver (if the Web service is acting as a client) uses the defined trust anchor to verify the signer certificate. The trust anchor is referenced using the trust anchor name.

To complete the signing information configuration process for request receiver, complete the following tasks:

1. Configure the server for request digital signature verification: Verifying the message parts
2. Configure the server for request digital signature verification: Choosing the verification method

To complete the process for the response receiver, if the Web services is acting as a client, complete the following tasks:

1. Configure the client for response digital signature verification: Verifying the message parts
2. Configure the client for response digital signature verification: Choosing the verification method

## Configuring trust anchors using the administrative console

This document describes how to configure trust anchors or trust stores at the application level. It does not describe how to configure trust anchors at the server or cell level. Trust anchors defined at the application level have a higher precedence over trust anchors defined at the server or cell level. For more information on creating and configuring trust anchors at the server or cell level, see either *Configuring the server security bindings using the Application Server Toolkit* or *Configuring the server security bindings using the administrative console*.

You can configure an application-level trust anchor using the Application Server Toolkit or the administrative console. This document describes how to configure the application-level trust anchor using the administrative console.

A trust anchor specifies key stores that contain trusted root certificates, which validate the signer certificate. These key stores are used by the request receiver (as defined in the `ibm-webservices-bnd.xmi` file) and the response receiver (as defined in the `ibm-webservicesclient-bnd.xmi` file when Web services is acting as client) to validate the signer certificate of the digital signature. The keystores are critical to the integrity of the digital signature validation. If they are tampered with, the result of the digital signature verification is doubtful and comprised. Therefore, it is recommended that you secure these keystores. The binding configuration specified for the request receiver in the `ibm-webservices-bnd.xmi` file must match the binding configuration for the response receiver in the `ibm-webservicesclient-bnd.xmi` file.

The steps in this document assume that you have already created a Web services-enabled Java 2 Platform, Enterprise Edition (J2EE) with Java Specification Requests (JSR) 109 enterprise application. If you have not created a Web services-enabled J2EE with JSR 109 enterprise application, see *Developing Web services*. Also, see either *Configuring the server security bindings using the Application Server Toolkit* or *Configuring the server security bindings using the administrative console* for an introduction on how to manage Web services security binding information for the server.

The following steps are for the client-side response receiver, which is defined in the `ibm-webservicesclient-bnd.xmi` file and the server-side request receiver, which is defined in the `ibm-webservices-bnd.xmi` file.

1. Click **Applications > Enterprise Applications > *enterprise\_application***.
2. In the Related Links section, click either **EJB Modules** or **Web Modules** and then click the Web services-enabled module in the **Uri** field.
3. Under Additional Properties, click **Web Services: Client Security Bindings** to edit the response receiver binding information, if Web services is acting as a client.
  - a. Under Response Receiver Binding, click **Edit**.
  - b. Under Additional Properties, click **Trust Anchors**.
  - c. Click **New** to create a new trust anchor.
  - d. Enter a unique name within the request receiver binding for the **Trust anchor name** field. The name is used to reference the trust anchor that is defined.
  - e. Enter the key store password, path, and key store type.
  - f. Click the trust anchor name link to edit the selected trust anchor.
  - g. Click **Remove** to remove the selected trust anchor or anchors.

When you start the application, the configuration is validated in the run time while the binding information is loading.
4. Return to the Web services-enabled module panel accessed in step 2.
5. Under Additional Properties, click **Web Services: Server Security Bindings** to edit the request receiver binding information.
  - a. Under Request Receiver Binding, click **Edit**.
  - b. Under Additional Properties, click **Trust Anchors**.
  - c. Click **New** to create a new trust anchor  
Enter a unique name within the request receiver binding for the **Trust anchor name** field. The name is used to reference the trust anchor that is defined.  
Enter the key store password, path, and key store type.  
Click the trust anchor name link to edit the selected trust anchor.  
Click **Remove** to remove the selected trust anchor or anchors.  
When you start the application, the configuration is validated in the run time while the binding information is loading.
6. Save the changes.

This procedure defines trust anchors that can be used by the request receiver or the response receiver (if the Web services is acting as client) to verify the signer certificate.

The request receiver or the response receiver (if the Web service is acting as a client) uses the defined trust anchor to verify the signer certificate. The trust anchor is referenced using the trust anchor name.

To complete the signing information configuration process for request receiver, complete the following tasks:

1. Configure the server for request digital signature verification: Verifying the message parts
2. Configure the server for request digital signature verification: Choosing the verification method

To complete the process for the response receiver, if the Web services is acting as client, complete the following tasks:

1. Configure the client for response digital signature verification: Verifying the message parts
2. Configure the client for response digital signature verification: Choosing the verification method

### ***Trust anchors collection:***



Use this page to view a list of keystore objects that contain trusted root certificates. These objects are used for certificate path validation of incoming X.509-formatted security tokens. Keystore objects within trust anchors contain trusted root certificates used by the CertPath API to validate the trust of a certificate chain.

To create the keystore file, use the key tool located in the `install_dir\java\jre\bin\keytool` directory.

To view this administrative console page, click **Servers > Application Servers > *server\_name***. Under Additional Properties, click **Web Services: Default bindings for Web Services Security > Trust Anchors**.

Click **New** to create a new trust anchor.

Click **Delete** to delete a trust anchor.

If you click **Update runtime**, the Web services security run time is updated with the default binding information, which is contained in the `ws-security.xml` file that was previously saved. If you make changes on this panel, you must complete the following steps:

1. Save your changes by clicking **Save** at the top of the administrative console. When you click **Save**, you are returned to the administrative console home panel.
2. Return to the Trust Anchors collection panel and click **Update runtime**. When you click **Update runtime**, the configuration changes made to the other Web services also are updated in the Web services security run time.

*Trust Anchor Name:*

Specifies the unique name used to identify the trust anchor.

*Key Store Path:*

Specifies the location of the keystore file that contains the trust anchors.

*Key Store Type:*

Specifies the type of keystore file.

The value for this field is either **JKS** or **JCEKS**.

#### ***Trust anchor configuration settings:***

Use this information to configure a trust anchor. Trust anchors point to key stores that contain trusted root or self-signed certificates. This information enables you to specify a name for the trust anchor and the information needed to access a key store. The application binding uses this name to reference a predefined trust anchor definition in the binding file (or the default).

To view this administrative console page, complete the following steps:

1. Click **Servers > Application Servers > *server\_name***.
2. Under Additional Properties, click **Web Services: Default bindings for Web Services Security > Trust Anchors > New**.

*Trust Anchor Name:*

Specifies the unique name used by the application binding to reference a predefined trust anchor definition in the default binding.

#### *Key Store Password:*

Specifies the password needed to access the key store file.

#### *Key Store Path:*

Specifies the location of the keystore file.

Use `${USER_INSTALL_ROOT}` as this path expands to the WebSphere Application Server path on your machine.

#### *Key Store Type:*

Specifies the type of key store file.

The value in this field is either **JKS** or **JCEKS**.

**JKS** Specify this option if you are not using Java Cryptography Extensions (JCE).

#### **JCEKS**

Specify this option if you are using Java Cryptography Extensions. Although the JCEKS key store format is more secure, it decreases performance.

<b>Data type</b>	String
<b>Default</b>	JKS
<b>Range</b>	JKS, JCEKS

### **Collection certificate store**

A *collection certificate store* is a collection of nonroot, certificate authority (CA) certificates and certificate revocation lists (CRLs). This collection of CA certificates and CRLs is used to check the signature of a digitally signed SOAP message.

The collection certificate stores are used when processing a received SOAP message. This collection is configured in the `securityRequestReceiverBindingConfig` section of the binding file for servers and in the `securityResponseReceiverBindingConfig` section of the binding file for clients.

A collection certificate store is one kind of certificate store. A certificate store is defined as `javax.security.cert.CertStore` in the Java CertPath application programming interface (API). The Java CertPath API defines the following types of certificate stores:

#### **Collection certificate store**

A collection certificate store accepts the certificates and CRLs as Java collection objects.

#### **Lightweight Directory Access Protocol certificate store**

The Lightweight Directory Access Protocol (LDAP) certificate store accepts certificates and CRLs as LDAP entries.

The CertPath API uses the certificate store and the trust anchor to validate the incoming X.509 certificate that is embedded in the SOAP message.

The Web services security implementation in the WebSphere Application Server supports the collection certificate store. Each certificate and CRL is passed as an encoded file. This configuration is done using either the administrative console or by scripting.

### **Configuring the client-side collection certificate store using the Assembly Toolkit**

A collection certificate store is a collection of non-root, certificate authority (CA) certificates and certificate revocation lists (CRLs). This collection of CA certificates and CRLs are used to check the signature of a digitally signed SOAP message.

You can configure the collection certificate either by using the Assembly Toolkit or the WebSphere Application Server administrative console. Complete the following steps to configure the client-side collection certificate store using the Assembly Toolkit.

1. Launch the Assembly Toolkit and either click **Windows > Open Perspective > J2EE**.
2. Select the Web services-enabled Enterprise JavaBean (EJB) or Web module.
3. In the Package Explorer window, locate the META-INF directory for an EJB module or the WEB-INF directory for a Web module.
4. Right-click the **webservicesclient.xml** file and select **Open With > Web Services Client Editor**.
5. Click the **Port Binding** tab in the Web Services Client Editor within the Assembly Toolkit. The Web Services Client Port Binding window is displayed.
6. Select one of the Port Qualified Name Binding entries.
7. Expand the **Security Response Receiver Binding Configuration > Certificate Store List > Collection Certificate Store** section.
8. Click **Add** to create a new collection certificate store, click **Edit** to edit an existing certificate store, or click **Remove** to delete an existing certificate store.
9. Enter a name in the **Name** field. This name is referenced in the **Certificate store reference** field in the Signing info dialog box.
10. Leave the **Provider** field as IBM CertPath.
11. Click **Add** to enter the path to your certificate store. For example, the path might be: `${USER_INSTALL_ROOT}/etc/ws-security/samples/intca2.cer`. If you have additional certificate store paths, click **Add** to add the paths.
12. Click **OK** when you finish adding paths.

## Configuring the client-side collection certificate store using the administrative console

A collection certificate store is a collection of non-root, certificate authority (CA) certificates and certificate revocation lists (CRLs). This collection of CA certificates and CRLs are used to check the signature of a digitally signed SOAP message.

You can configure the collection certificate either by using the WebSphere Application Server Toolkit or the WebSphere Application Server administrative console. Complete the following steps to configure the client-side collection certificate store using the administrative console.

1. Connect to the WebSphere Application Server administrative console. You can connect to the administrative console by typing `http://localhost:9090/admin` in your Web browser unless you have changed the port number.
2. Click **Applications > Enterprise Applications > *application\_name***.
3. Under Related Items, click either **Web Modules** or **EJB Modules** depending on the type of module you are securing.
4. Click the name of the module you are securing.
5. Under Additional Properties, click either **Web Services: Client Security Bindings** to add the collection certificate store to the client security bindings. If you do not see any entries, return to the WebSphere Application Server Toolkit and configure the security extensions for either the client or the server.

To configure the security extensions for the client, see the following topics:

- Configuring the client for response digital signature verification: Verifying the message parts
  - Configuring the client for response digital signature verification: Choosing the verification method
6. Under Response Receiver Binding, click **Edit** to edit the client security bindings.
  7. Click **Collection Certificate Store**.

8. Click a Certificate Store Name to edit an existing certificate store or click **New** to add a new certificate store name.
9. Enter a name in the **Certificate Store Name** field. The name entered in this field is a name that is referenced in the **Certificate Store** field on the Signing information configuration page.
10. Leave the **Certificate Store Provider** field value as IBM CertPath.
11. Click **Apply**.
12. Under Additional Properties, click **X.509 Certificates > New**.
13. Enter the path to your certificate store. For example, the path might be: `${USER_INSTALL_ROOT}/etc/ws-security/samples/intca2.cer`. If you have any additional certificate store paths to enter, click **New** and add the path names.
14. Click **OK**.

## Configuring the server-side collection certificate store using the Assembly Toolkit

A collection certificate store is a collection of non-root, certificate authority (CA) certificates and certificate revocation lists (CRLs). This collection of CA certificates and CRLs are used to check the signature of a digitally signed SOAP message.

You can configure the collection certificate either by using the Assembly Toolkit or the WebSphere Application Server administrative console. Complete the following steps to configure the server-side collection certificate store using the Assembly Toolkit.

1. Launch the Assembly Toolkit and click **Windows > Open Perspective > J2EE**.
2. Select the Web services-enabled Enterprise JavaBean (EJB) or Web module.
3. In the Package Navigator window, locate the META-INF directory for an EJB module or the WEB-INF directory for a Web module.
4. Right-click the `webservices.xml` file, select **Open With > Web Services Editor**.
5. Click the **Binding Configurations** tab in the Web services editor within the Assembly Toolkit. The Web Service Binding Configuration window is displayed.
6. Select one of the Web service description binding entries under the Port Component Binding section.
7. Expand the **Request Receiver Binding Configuration Details > Certificate Store List > Collection Certificate Store** section.
8. Click **Add** to create a new collection certificate store, click **Edit** to edit an existing certificate store, or click **Remove** to delete an existing certification store.
9. Enter a name in the **Name** field. This name is referenced in the **Certificate store reference** field in the Signing info dialog.
10. Leave the **Provider** field as IBM CertPath.
11. Click **Add** to enter the path to your certificate store. For example, the path might be: `${USER_INSTALL_ROOT}/etc/ws-security/samples/intca2.cer`. If you have additional certificate store paths, click **Add** to add the paths.
12. Click **OK** when you finish adding paths.

### **Collection certificate store collection:**

Use this page to view a list of certificate stores containing untrusted, intermediary certificate files awaiting validation. Validation might consist of checking to see if the certificate is on a certificate revocation list (CRL), checking that the certificate has not expired, and checking that the certificate was issued by a trusted signer.

To view this administrative console page, complete the following steps:

1. Click **Servers > Application Servers > server\_name**.

2. Under Related Items, click **Web Services: Default bindings for Web Services Security > Collection Certificate Store**.
3. Click **New** to specify a store name and provider for a new collection certificate store.
4. Click **Delete** to delete a collection certificate store.

Using this panel, complete the following steps:

1. Specify a certificate store name and certificate store provider.
2. Save your changes by clicking **Save** at the top of the administrative console. When you click **Save**, you return to the administrative console home panel.
3. Return to the collection certificate store collection panel and click **Update runtime** to update the Web services security run time with the default binding information, which is found in the `ws_security.xml` file.

When you click **Update runtime**, the configuration changes made to the other Web services also are updated in the Web services security run time.

*Certificate Store Name:*

Specifies the name of the certificate store.

*Certificate Store Provider:*

Specifies the provider of the certificate store.

#### **Collection certificate store configuration settings:**

Use this page to specify the name and provider of a certificate store.

To view this administrative console page, complete the following steps:

1. Click **Servers > Application Servers > *server\_name***.
2. Under **Additional Properties**, click **Web Services: Default bindings for Web Services Security > Collection Certificate Store > New**.

*Certificate Store Name:*

Specifies the name for the certificate store. The application binding uses the certificate store name to reference a predefined binding.

*Certificate Store Provider:*

Specifies the provider for the certificate store implementation.

<b>Data type</b>	String
<b>Default</b>	IBM CertPath

#### **X.509 certificates collection:**

Use this page to view a list of X.509 certificates.

To view this administrative console page, complete the following steps:

1. Click **Servers > Application Servers > *server\_name***.
2. Under **Additional Properties**, click **Web Services: Default bindings for Web Services Security > Collection Certificate Store**.

3. On the Collection Certificate Store page under Additional Properties, click **X.509 Certificates**.
4. Click **New** to create a new path to an X.509 certificate or click **Delete** to delete a path to an X.509 certificate.

*X509 Certificate Path:*

Specifies the location of the X.509 certificate.

#### ***X.509 certificate configuration settings:***

Use this page to specify the location of your X.509 certificates.

To view this administrative console page, complete the following steps:

1. Click **Servers > Application Servers > *server\_name***.
2. Under **Additional Properties**, click **Web Services: Default bindings for Web Services Security > Collection Certificate Store**.
3. On the Collection Certificate Store page, under **Additional Properties**, click **X.509 Certificates > New**.

*X509 Certificate Path:*

Specifies the location of the X.509 certificate.

## **Configuring the server-side collection certificate store using the administrative console**

A *collection certificate store* is a collection of nonroot, certificate authority (CA) certificates and certificate revocation lists (CRLs). This collection of CA certificates and CRLs is used to check the signature of a digitally signed SOAP message.

You can configure the collection certificate either by using the WebSphere Application Server Toolkit or the WebSphere Application Server administrative console. Complete the following steps to configure the server-side collection certificate store using the administrative console.

1. Connect to the WebSphere Application Server administrative console. You can connect to the administrative console by typing `http://localhost:9090/admin` in your Web browser unless you have changed the port number.
2. Click **Applications > Enterprise Applications > *application\_name***.
3. Under Related Items, click either **Web Modules** or **EJB Modules** depending on the type of module you are securing.
4. Click the name of the module you are securing.
5. Under Additional Properties, click **Web Services: Server Security Bindings** to add the collection certificate store to the server security bindings. If you do not see any entries, return to the WebSphere Application Server Toolkit and configure the security extensions for the server.

To configure the security extensions for the server, see the following topics:

- Configuring the server for request digital signature verification: verifying the message parts
  - Configuring the server for request digital signature verification: choosing the verification method
6. Click **Edit** under Request Receiver Binding to edit the server security bindings.
  7. Click **Collection Certificate Store**.
  8. Click a Certificate Store Name to edit an existing certificate store or click **New** to add a new certificate store name.
  9. Enter a name in the **Certificate Store Name** field. The name entered in this field is a name that is referenced in the **Certificate Store** field on the Signing information configuration page.

10. Leave the **Certificate Store Provider** field as IBM CertPath.
11. Click **Apply**.
12. Under Additional Properties, click **X.509 Certificates > New**.
13. Enter the path to your certificate store. For example, the path might be: `${USER_INSTALL_ROOT}/etc/ws-security/samples/intca2.cer`. If you have any additional certificate store paths to enter, click **New** and add the path names.
14. Click **OK**.

## Configuring default collection certificate stores at the server level in the WebSphere Application Server administrative console

A collection certificate store is a collection of non-root, certificate authority (CA) certificates and certificate revocation lists (CRLs). This collection of CA certificates and CRLs are used to check the signature of a digitally signed SOAP message. A certificate store typically refers to a certificate store located in the file system. The location of the certificate store can vary from machine to machine, so you might configure a default collection certificate store for a specific machine and reference it from within the signing information. The signing information is found within the binding configurations of any application installed on the machine. This suggestion enables you to define a single collection certificate store for all of the applications that need to use the same certificates. You also can specify the default binding information at the cell level.

Complete the following steps to configure the default collection certificate store at the server level using the WebSphere Application Server administrative console:

1. Connect to the administrative console. You can access the administrative console by typing `http://localhost:9090/admin` in your Web browser unless you have changed the port number.
2. Click **Servers > Application Servers > server1**.
3. Under Additional Properties, click **Web Services: Default bindings for Web Services Security > Collection Certificate Store**.
4. Enter a name in the **Certificate Store Name** field. This name is referenced in the Certificate Store field on the Signing information configuration page.
5. Leave the **Certificate Store Provider** field value as IBM CertPath.
6. Click **Apply**.
7. Under Additional Properties, click **X.509 Certificates > New**.
8. Enter the path to your certificate store. For example, the path might be: `${USER_INSTALL_ROOT}/etc/ws-security/samples/intca2.cer`.  
If you have any additional certificate store paths to enter, click **New** and add the path names.
9. Click **OK**.

### Key locator

A *key locator* (`com.ibm.wsspi.wssecurity.config.KeyLocator`) is an abstraction of the mechanism that retrieves the key for digital signature and encryption.

You can use any of the following infrastructure from which to retrieve the keys depending upon the implementation:

- Java keystore file
- Database
- Lightweight Third Party Authentication (LDAP) server

Key locators search the key using some type of a clue. The following types of clues are supported:

- A string label of the key, which is explicitly passed through the application programming interface (API). The relationships between each key and its name (string label) is maintained inside the key locator.



- The execution context of the key locator; explicit information is not passed to the key locator. A key locator determines the appropriate key according to the execution context.

Current versions of key locators do not support the retrieval of verification keys because current Web services security implementations do not support the secret key-based signature. Because the key locators support the public key-based signature only, the key for verification is embedded in the X.509 certificate as a <BinarySecurityToken> element in the incoming message.

For example, key locators can obtain the identity of the caller from the context and can retrieve the public key of the caller for response encryption.

## Usage scenarios

This section describes the usage scenarios for key locators.

### Signing

The name of the signing key is specified in the Web services security configuration. This value is passed to the key locator and the actual key is returned. The corresponding X.509 certificate also can be returned.

### Verification

As described previously, key locators are not used in signature verification.

### Encryption

The name of the encryption key is specified in the Web services security configuration. This value is passed to the key locator and the actual key is returned.

### Decryption

The Web services security specification recommends using the key identifier instead of the key name. However, while the algorithm for computing the identifier for the public keys is defined in Internet Engineering Task Force (IETF) Request for Comment (RFC) 3280, there is no agreed upon algorithm for the secret keys. Therefore, the current implementation of Web services security uses the identifier only when public key-based encryption is performed. Otherwise, the ordinal key name is used.

When you use public key-based encryption, the value of the key identifier is embedded in the incoming encrypted message. Then, the Web services security implementation searches for all the keys managed by the key locator and decrypts the message using the key whose identifier value matches the one in the message.

When you use secret key-based encryption, the value of the key name is embedded in the incoming encrypted message. The Web services security implementation asks the key locator for the key with the name that matches the name in the message and decrypts the message using the key.

### ***Key locator collection:***

Use this page to view a list of available key locators. Key locators identify the keys needed for digital signature and encryption. A key locator must implement the `com.ibm.wsspi.wssecurity.config.KeyLocator` interface. The two default implementations are:

`com.ibm.wsspi.wssecurity.config.WSIdKeyStoreMapKeyLocator` and  
`com.ibm.wsspi.wssecurity.config.KeyStoreKeyLocator`.

To view this administrative console page, complete the following steps:

1. Click **Servers > Application Servers > *server\_name***.

2. Under Additional Properties, click **Web Services: Default bindings for Web Services Security > Key Locators**.
3. Click **New** to create a key locator. Click **Delete** to delete a key locator.

Using this **Key Locator collection** panel, complete the following steps:

1. Specify a key locator name and key locator class name on the panel
2. Save your changes by clicking **Save** at the top of the administrative console. The administrative console home panel is displayed.
3. After saving your changes, return to the **Key Locator collection** panel to update the Web services security run time with the default binding information, which is found in the `ws_security.xml` file.
4. To update the Web services security run time, click **Update runtime**. When you click **Update runtime**, the configuration changes made to the other Web services also are updated in the Web services security run time.
5. After you define key locators, click the key locator name to specify additional properties and keys under **Additional Properties**.

*Key Locator Name:*

Specifies the unique name of the key locator.

*Key Locator Classname:*

Specifies the class name of the key locator in the key store file.

#### **Key locator configuration settings:**

Use this page to specify the settings for key locators.

To view this administrative console page, complete the following steps:

1. Click **Servers > Application Servers > server\_name**.
2. Under Additional Properties, click **Web Services: Default bindings for Web Services Security > Key Locators > New**.

*Key Locator Name:*

Specifies the name of the key locator.

**Data type** String

*Key Locator Classname:*

Specifies the name for the key locator class implementation.

WebSphere Application Server has the following default key locator class implementations:

#### **com.ibm.wsspi.wssecurity.config.WSIdKeyStoreMapKeyLocator**

Maps an authenticated identity to a key. This class is used by the response sender. If encryption is used, this class is used to locate a key to encrypt the response message. The `com.ibm.wsspi.wssecurity.config.WSIdKeyStoreMapKeyLocator` class has the capability to map an authenticated identity from the invocation credential of the current thread to a key that is used to encrypt the message. If an authenticated identity is present on the current thread, the class maps the ID to the mapped name. For example, `user1` is mapped to `mappedName_1`. Otherwise, `name="default"`. When a matching key is not found, the authenticated identity is mapped to the default key specified in the binding file.

## **com.ibm.wsspi.wssecurity.config.KeyStoreKeyLocator**

Maps a name to an alias. This class is used by the response receiver, request sender, and request receiver. The encryption process uses this class to obtain a key to encrypt a message, and the digital signature process uses this class to obtain a key to sign a message. The `com.ibm.wsspi.wssecurity.config.KeyStoreKeyLocator` class maps a logical name to a key alias in the keystore file. For example, key #105115176771 is mapped to CN=Alice, O=IBM, c=US.

<b>Data type</b>	String
<b>Defaults</b>	<code>com.ibm.wsspi.wssecurity.config.KeyStoreKeyLocator</code>  <code>com.ibm.wsspi.wssecurity.config.WSIdKeyStoreMapKeyLocator</code>

### *Key Store Password:*

Specifies the password used to access the keystore file.

### *Key Store Path:*

Specifies the location of the keystore file.

Use `${USER_INSTALL_ROOT}` as this path expands to the WebSphere Application Server path on your machine.

### *Key Store Type:*

Specifies the type of keystore file.

The value for this field is either JKS or JCEKS:

**JKS** Use this option if you are not using Java Cryptography Extensions (JCE).

#### **JCEKS**

Use this option if you are using Java Cryptography Extensions.

<b>Default</b>	JKS
<b>Range</b>	JKS, JCEKS

## **Keys**

Keys are used for XML signature and encryption.

There are two predominant kinds of keys used in the current Web services security implementation:

- Public key - such as Rivest Shamir Adleman (RSA) encryption and Digital Signature Algorithm (DSA) encryption
- Secret key - such as Data Encryption Standard (DES) encryption

In public key-based signature, a message is signed using sender private key and is verified using the sender public key. In public key-based encryption, a message is encrypted using the receiver public key and is decrypted using the receiver private key. In secret key-based signature and encryption, the same key is used by both parties.

While the current implementation of Web services security can support both kinds of keys, there are a few items to note:

- Secret key-based signature is not supported.
- The format of the message differs slightly between public key-based encryption and secret key-based encryption.

### **Key collection:**

Use this page to view a list of logical names that are mapped to a key alias in the key store file.

To view this administrative console page, complete the following steps:

1. Click **Servers > Application Servers > *server\_name***.
2. Under Additional Properties, click **Web Services: Default bindings for Web Services Security > Key Locators > *key\_locator\_name***.
3. Under Additional Properties, click **Keys**.
4. Click **New** to create a new key object in the key store file.
5. Click **Delete** to delete a mapping of a key object within the key store file.

*Key Name:*

Specifies the name of the key object found in the key store file.

*Key Alias:*

Specifies an alias for the key object.

The alias is used when the key locator searches for the key objects in the key store.

### **Key configuration settings:**

Use this page to define the mapping of a logical name to a key alias in a key store file.

To view this administrative console page, complete the following steps:

1. Click **Servers > Application Servers > *server\_name***.
2. Under Additional Properties, click **Web Services: Default bindings for Web Services Security > Key Locators > *key\_locator\_name***.
3. Under Additional Properties, click **Keys > New**.

*Key Name:*

Specifies the name of the key object. This name is used by the key locator to find the key within the key store file.

*Key Alias:*

Specifies the alias for the key object contained in the key store file.

*Key Password:*

Specifies the password needed to access the key object within the key store file.

## **Web services security service provider programming interfaces**

Several Service Provider Interfaces (SPIs) are provided to extend the capability of the Web services security run time. The following list contains the SPIs that are available for WebSphere Application Server:

- `com.ibm.wsspi.wssecurity.config.KeyLocator` is an abstract for obtaining the keys for digital signature and encryption. The following list contains the default implementations:
  - `com.ibm.wsspi.wssecurity.config.KeyStoreKeyLocator`  
Implements the Java key store.
  - `com.ibm.wsspi.wssecurity.config.WSIdKeyStoreMapKeyLocator`

Provides a mapping of the authenticated identity to a key for encryption. Or, the implementation uses the default key that is specified. This implementation is typically used in the response sender configuration.

- `com.ibm.wsspi.wssecurity.config.CertInRequestKeyLocator`

Provides the capability of using the signer key for encryption in the response message. This implementation is typically used in the response sender configuration.

- `com.ibm.wsspi.wssecurity.id.TrustedIDEvaluator` is an interface that is used to evaluate the trust for identity assertion. The default implementation is `com.ibm.wsspi.wssecurity.id.TrustedIDEvaluatorImpl`, which enables you to define a list of trusted identities.
- The Java Authentication and Authorization Service (JAAS) CallbackHandler application programming interfaces (APIs) are used for token generation by the request sender. This interface can be extended to generate a custom token that can be inserted in the Web services security header. The following list contains the default implementations that are provided by WebSphere Application Server:
  - `com.ibm.wsspi.wssecurity.auth.callback.GUIPromptCallbackHandler`  
Presents a login prompt to gather the basic authentication data. Use this implementation in the client environment only.
  - `com.ibm.wsspi.wssecurity.auth.callback.StdinPromptCallbackHandler`  
Collects the basic authentication data in the standard in (stdin) prompt. Use this implementation in the client environment only.
  - `com.ibm.wsspi.wssecurity.auth.callback.NonPromptCallbackHandler`  
Reads the basic authentication data from the application binding file. This implementation might be used on the server side to generate a user name token.
  - `com.ibm.wsspi.wssecurity.auth.callback.LTPATokenCallbackHandler`  
Generates a Lightweight Third Party Authentication (LTPA) token in the Web services security header as a binary security token. If basic authentication data is defined in the application binding file, it is used to perform a login, to extract the LTPA token from the WebSphere credentials, and to insert the token in the Web services security header. Otherwise, it will extract the LTPA security token from the invocation credentials (run as identity) and insert the token in the Web services security header.

The JAAS LoginModule API is used for token validation on the request receiver side of the message. You can implement a custom LoginModule API to perform validation of the custom token on the request receiver of the message. After the token is verified and validated, the token is set as the caller and then run as the identity in the WebSphere Application Server run time. The identity is used for authorization checks by the containers before a Java 2 Platform, Enterprise Edition (J2EE) resource is invoked. The following list presents the are the default AuthMethod configurations provided by WebSphere Application Server:

#### **BasicAuth**

Validates a user name token.

#### **Signature**

Maps the distinguished name (DN) of a verified certificate to a Java Authentication and Authorization Service (JAAS) subject.

#### **IDAssertion**

Maps a trusted identity to a JAAS subject.

**LTPA** Validates an LTPA token that is received in the message and creates a JAAS subject.

## **Configuring key locators using the Assembly Toolkit**

This task provides instructions on how to configure key locators using the Assembly Toolkit. You can configure key locators in various locations within the Assembly Toolkit. This task provides instructions on how to configure key locators at any of these locations because the concept is the same.

1. Launch the Assembly Toolkit and click **Windows > Open Perspective > J2EE**.
2. Select the Web services-enabled Enterprise JavaBean (EJB) or Web module.
3. In the Package Explorer window, locate the META-INF directory for an EJB module or the WEB-INF directory for a Web module.
4. Right-click the **webservicesclient.xml** file and click **Open With > Web Services Client Editor** or right-click the **webservices.xml** file and click **Open With > Web Services Editor**.
5. Click the **Port Binding** tab in the Web Services Client Editor within the Assembly Toolkit or the **Binding Configurations** tab in the Web Services Editor within the Assembly Toolkit.
6. Expand one of the **Binding Configuration** sections.
7. Expand the **Key Locators** section.
8. Click **Add** to create a new key locator, click **Edit** to edit an existing key locator, or click **Remove** to delete an existing key locator.
9. Enter a key locator name. The name entered for the **Key locator name** is used to refer to the key locator from the Encryption information and Signing Information sections.
10. Enter a key locator class. The key locator class is the implementation of the KeyLocator interface. When using default implementations, select a class from the menu.
11. Determine whether to click **Use key store**. Select this option when you use the default implementations as they use key stores. If you click **Use key store**, complete the following steps:
  - a. Enter a value in the **key store storepass** field. The key store storepass is the password used to access the key store.
  - b. Enter a path name in the **key store path** field. The key store path is the location on the file system where the key store resides. Make sure that the location can be found wherever you deploy the application.
  - c. Enter a type value in the **key store type** field. The valid types to enter are JKS and JCEKS. JKS is used when you are not using Java Cryptography Extensions (JCE). JCEKS is used when you are using JCE. Although the JCEKS type is more secure, it might decrease performance.
  - d. Click **Add** to create an entry for a key in the key store.
    - 1) Enter a value in the **Alias** field.  
The key alias is a reference to this particular key from the Signing Information section.
    - 2) Enter a value in the **Key pass** field.  
The key pass is the password associated with the certificate which is created using the Development Kit, Java Technology Edition keytool.exe file.
    - 3) Enter a value in the **Key name** field.  
The key name refers to the alias of the certificate as found in the key store.
12. Click **Add** to create a custom property. The property can be used by custom key locator implementations. For example, you can use properties with the WsIdKeyStoreMapKeyLocator default implementation. The key locator implementation has the following property names:
  - *id\_*, which maps to a credential user ID.
  - *mappedName\_*, which maps to the key alias to use for this user name.
  - *default*, which maps to a key alias to use when a credential does not have an associated *id\_* entry.

A typical set of properties for this key locator might be: `id_1=user1, mappedName_1=key1, id_2=user2, mappedName_2=key2, default=key3`. If user1 or user2 authenticates, then the associated key1 or key2 is used, respectively. However, if none of the user properties authenticate or the user is not user1 or user2, then key3 is used.

  - a. Enter a name in the **Name** field. The name entered is the property name.
  - b. Enter a value in the **Value** field. This value entered is the property value.



## Configuring key locators using the administrative console

This task provides instructions on how to configure key locators using the WebSphere Application Server administrative console. You can configure binding information in the administrative console. You must use the Assembly Toolkit to configure extensions. The following steps are used to configure a key locator in the administrative console for a specific application:

1. Connect to the administrative console by typing `http://:9090/admin` in your Web browser unless you have changed the port number.
2. Click **Applications > Enterprise Applications > *application\_name***.
3. Under Related Items, click either **Web Modules** or **EJB Modules**, depending on the type of module you are securing.
4. Click the name of the module you are securing.
5. Under Additional Properties, click either **Web Services: Client Security Bindings** or **Web Services: Server Security Bindings**, depending on whether you are adding the key locator to the client security bindings or to the server security bindings. If you do not see any entries, return to the Assembly Toolkit and configure the security extensions.
6. Edit the Request Sender Binding, Response Receiver Binding, Request Receiver Binding, or Response Sender Binding.
  - If you are editing your client security bindings, click **Edit** for either the Request Sender Binding or the Response Receiver Binding.
  - If you are editing your server security bindings, click **Edit** for either the Request Receiver Binding or the Response Sender Binding.
7. Click **Key Locators**.
8. Click **New** to configure a new key locator, select the box next to a key locator name and click **Delete** to delete a key locator, or click the name of a key locator to edit its configuration. If you are configuring a new key locator or editing an existing one, complete the following steps:
  - a. Specify a name for the key locator in the **Key Locator Name** field.
  - b. Specify a name for the key locator class implementation in the **Key Locator Classname** field. WebSphere Application Server has the following default key locator class implementations:

### **com.ibm.wsspi.wssecurity.config.WSIdKeyStoreMapKeyLocator**

This class is used by the response sender to map an authenticated identity to a key. If encryption is used, this class is used to locate a key to encrypt the response message. The `com.ibm.wsspi.wssecurity.config.WSIdKeyStoreMapKeyLocator` class has the capability to map an authenticated identity from the invocation credential of the current thread to a key that is used to encrypt the message. If an authenticated identity is present on the current thread, the class maps the ID to the mapped name. For example, `user1` is mapped to `mappedName_1`. Otherwise, `name="default"`. When a matching key is not found, the authenticated identity is mapped to the default key specified in the binding file.

### **com.ibm.wsspi.wssecurity.config.KeyStoreKeyLocator**

This class is used by the response receiver, the request sender, and the request receiver to map a name to an alias. Encryption uses this class to obtain a key to encrypt a message and digital signature uses this class to obtain a key to sign a message. The `com.ibm.wsspi.wssecurity.config.KeyStoreKeyLocator` class maps a logical name to a key alias in the key store file. For example, `key #105115176771` maps to `CN=Alice, O=IBM, C=US`.

- c. Specify the password used to access the key store password in the **Key Store Password** field. This field is optional because the key locator does not use a key store.
- d. Specify the path name used to access the key store in the **Key Store Path** field. This field is optional because the key locator does not use a key store. Use `${USER_INSTALL_ROOT}` because this path expands to the WebSphere Application Server path on your machine.



- e. Select a keystore type from the **Key Store Type** field. This field is optional because the key locator does not use a key store. Use the **JKS** option if you are not using Java Cryptography Extensions (JCE) and use **JCEKS** if you are using JCE.

### Trusted ID evaluator

Trusted ID evaluator (com.ibm.wsspi.wssecurity.id.TrustedIDEvaluator) is an abstraction of the mechanism that evaluates whether the given ID name is trusted.

Depending upon the implementation, you can use various types of infrastructure to store a list of the trusted IDs, such as:

- Plain text file
- Database
- Lightweight Directory Access Protocol (LDAP) server

The trusted ID evaluator is typically used by the eventual receiver in a multi-hop environment. The Web services security implementation invokes the trusted ID evaluator and passes the identity name of the intermediary as a parameter. If the identity is evaluated and deemed trustworthy, the procedure continues. Otherwise, an exception is thrown and the procedure is stopped.

#### **Trusted ID evaluator collection:**

Use this page to view a list of trusted identity (ID) evaluators. The trusted ID evaluator determines whether to trust the identity-asserting authority. Once the ID is trusted, the WebSphere Application Server issues the proper credentials, which are used in a downstream call for invoking resources. The trusted ID evaluator implements the com.ibm.wsspi.wssecurity.id.TrustedIDEvaluator interface.

To view this administrative console page, complete the following steps:

1. Click **Servers > Application Servers > server\_name**.
2. Under Additional Services, click **Web Services: Default bindings for Web Services Security > Trusted ID Evaluators**.
3. Click **New** to create a trusted ID evaluator or click **Delete** to delete a trusted ID evaluator.

Using this trusted ID evaluator collection panel, complete the following steps:

1. Specify a trusted ID evaluator name and trusted ID evaluator class name.
2. Save your changes by clicking **Save** at the top of the administrative console. When you click **Save**, you return to the administrative console home panel.
3. Return to the Trusted ID Evaluator collection panel to update the Web services security run time with the default binding information, which is found in the ws\_security.xml file.
4. Click **Update runtime**. The configuration changes made to the other Web services also are updated in the Web services security run time.

#### *Trusted ID Evaluator Name:*

Specifies the unique name of the trusted ID evaluator.

#### *Trusted ID Evaluator Classname:*

Specifies the class name of the trusted ID evaluator.

#### **Trusted ID evaluator configuration settings:**

Use this information to configure trust identity (ID) evaluators.

To view this administrative console page, complete the following steps:

1. Click **Servers > Application Servers > server\_name**.
2. Under Additional Properties, click **Web Services: Default bindings for Web Services Security > Trust ID Evaluators > New**.

You must specify the name and value properties for the default trusted ID evaluator to create the trusted ID list for evaluation.

*Trusted ID Evaluator Name:*

Specifies the unique name used by the application binding to refer to a trusted identity (ID) evaluator defined in the default binding.

You must specify the trusted ID evaluator name in the form, `trustedId_n`, where `n` is an integer from 0 to `n`.

*Trusted ID Evaluator Class Name:*

Specifies the class name of the trusted ID evaluator.

**Default**

`com.ibm.wsspi.wssecurity.id.TrustedIDEvaluatorImpl`

## Login mappings

Login mappings, found in the `ibm-webservices-bnd.xml` Extended Markup Language (XML) file, contains a mapping configuration. This mapping configuration defines how the Web services security handler maps the token `<ValueType>` element that is contained within the security token extracted from the message header, to the corresponding authentication method. The token `<ValueType>` element is contained within the security token extracted from a SOAP message header.

The sender-side Web services security handler generates and attaches security tokens based on the `<AuthMethods>` element that is specified in the deployment descriptor. For example, if the authentication method is `BasicAuth`, the sender-side security handler generates and attaches `UsernameToken` (with both user name and password) to the SOAP message header. The Web services security run time uses the Java Authentication and Authorization Service (JAAS) `javax.security.auth.callback.CallbackHandler` interface as a security provider to generate security tokens on the client side (or when Web services is acting as client).

The sender security handler invokes the `handle()` method of a `javax.security.auth.callback.CallbackHandler` interface implementation. This implementation creates the security token and passes the token back to the sender security handler. The sender's security handler constructs the security token based on the authentication information in the callback array. The security handler then inserts the security token into the Web Services Security message header.

The `CallbackHandler` interface implementation that you use to generate the required security token is defined in the `<loginBinding>` element in the `ibm-webservicesclient-bnd.xml` Web services security binding file. For example,

```
<loginBinding xmi:id="LoginBinding_1052760331526" authMethod="BasicAuth"
 callbackHandler="com.ibm.wsspi.wssecurity.auth.callback.StdinPromptCallbackHandler"/>
```

The `<loginBinding>` element associates the `com.ibm.wsspi.wssecurity.auth.callback.StdinPromptCallbackHandler` interface with the `BasicAuth` authentication method. WebSphere Application Server provides the following set of `CallbackHandler` interface implementations you can use to create various security token types:

### **com.ibm.wsspi.wssecurity.auth.callback.GUIPromptCallbackHandler**

If there is no basic authentication data defined in the login binding information (this information is not the same as the HTTP basic authentication information), the previous token type prompts for

user name and password through a login panel. The implementation uses the basic authentication data defined in the login binding. Use this CallbackHandler with the BasicAuth authentication method. Do not use this CallbackHandler implementation on the server because it prompts you for login binding information.

**com.ibm.wsspi.wssecurity.auth.callback.StdinPromptCallbackHandler**

If basic authentication data is not defined in the login binding (this information is not the same as the HTTP basic authentication information), the implementation prompts for the user name and password using standard in (stdin). The implementation uses the basic authentication data defined in the login binding. Use this CallbackHandler implementation with the BasicAuth authentication method. Do not use this CallbackHandler implementation on the server because it prompts you for login binding information.

**com.ibm.wsspi.wssecurity.auth.callback.NonPromptCallbackHandler**

This CallbackHandler implementation does not prompt. Rather, it uses the basic authentication data defined in the login binding (this information is not the same as the HTTP basic authentication information). This CallbackHandler implementation is meant for use with the BasicAuth authentication method. You must define the basic authentication data in the login binding information for this CallbackHandler implementation. You can use this implementation when Web services is running as a client and needs to send basic authentication (<wsse:UsernameToken>) to the downstream call.

**com.ibm.wsspi.wssecurity.auth.callback.LTPATokenCallbackHandler**

The CallbackHandler generates Lightweight Third Party Authentication (LTPA) tokens from the run as JAAS Subject (invocation Subject) of the current WebSphere Application Server security context. However, if basic authentication data is defined in the login binding information (not the HTTP basic authentication information), the implementation uses the basic authentication data and LTPA token generated. The **Token Type URI** and **Token Type Local Name** values must be defined in the login binding information for this CallbackHandler implementation. The token value type is used to validate the token to the request sender and request receiver binding configuration. The Web services security run time inserts the LTPA token as a binary security token (<wsse:BinarySecurityToken>) into the message SOAP header. The value type is mandatory. (See LTPA for more information). Use this CallbackHandler implementation with the LTPA authentication method.

Figure 1 shows the sender security handler in the request sender message process.

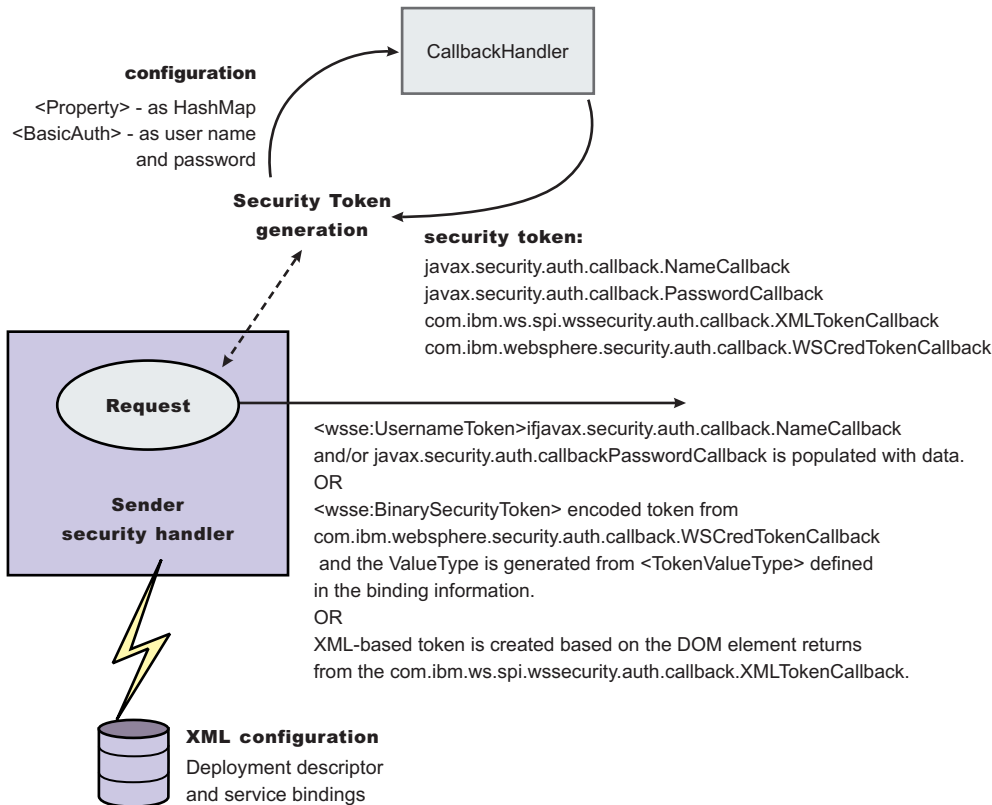


Figure 23. Request sender SOAP message process

You can configure the receiver-side security server to support multiple authentication methods and multiple types of security tokens. The following steps describe the request sender SOAP message process:

1. After receiving a message, the receiver Web services security handler compares the token type (in the message header) with the expected token types configured in the deployment descriptor.
2. The Web services security handler extracts the security token form the message header and maps the token <ValueType> element to the corresponding authentication method. The mapping configuration is defined in the <loginMappings> element in the `ibm-webservices-bnd.xml` XML file. For example:

```
<loginMappings xmi:id="LoginMapping_1051977980074" authMethod="LTPA"
 configName="WSLogin">
 <callbackHandlerFactory xmi:id="CallbackHandlerFactory_1051977980081"
 classname="com.ibm.wsspi.wssecurity.auth.callback.WSCallbackHandlerFactoryImpl"/>
 <tokenValueType xmi:id="TokenValueType_1051977980081"
 uri="http://www.ibm.com/websphere/appserver/tokenype/5.0.2" localName="LTPA"/>
</loginMappings>
```

The `com.ibm.wsspi.wssecurity.auth.callback.CallbackHandlerFactory` interface is a factory for `javax.security.auth.callback.CallbackHandler`.

3. The Web services security run time initiates the factory implementation class and passes the authentication information from Web services security header to the factory class through the set methods.
4. The Web services security run time invokes the `newCallbackHandler()` method to obtain the `javax.security.auth.CallbackHandler` object, which generates the required security token.
5. When the security handler receives an LTPA `BinarySecurityToken`, it uses the `WSLogin` JAAS login configuration and the `newCallbackHandler()` method to validate the security token. If none of the expected token types are found in the SOAP message Web services security header, the request is rejected with an SOAP fault. Otherwise, the token type is used to map to a JAAS login configuration

for token validation. If authentication is successful, a JAAS Subject is created and associated with the running thread. Otherwise, the request is rejected with a SOAP fault.

The following table shows the authentication methods and the JAAS login configurations.

Authentication method	JAAS login configuration
BasicAuth	WSLogin
Signature	system.wssecurity.Signature
LTPA	WSLogin
IDAssertion	system.wssecurity.IDAssertion

Figure 2 shows the receiver security handler in the request receiver message process.

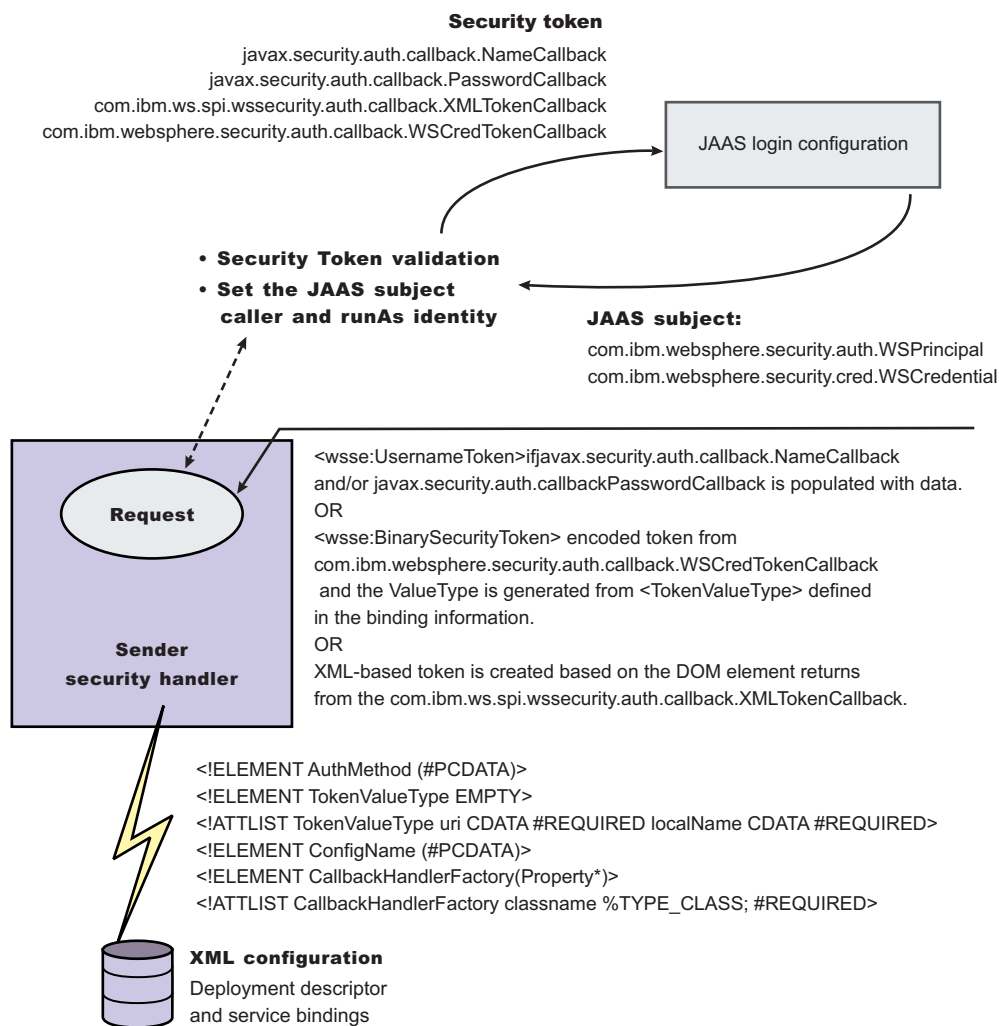


Figure 24. Request receiver SOAP message process

The default <LoginMapping> is defined in the server-level ws-security.xml files. If nothing is defined in the binding file information, the ws-security.xml default is used. However, an administrator can override the default by defining a new <LoginMapping> element in the binding file.

6. The client reads the default binding information in the `${install_dir}/properties/ws-security.xml` file.
7. The server run-time component loads the server-level `ws-security.xml` files, if they exist. The two files are merged in the run time to form one effective set of default binding information.

On a base application server, the server run time component only loads the server-level `ws-security.xml` file. The server-side `ws-security.xml` file and the application Web services security binding information are managed by the administrative console and by the WSADMIN. You can specify the binding information during application deployment either through the administrative console or through the WSADMIN file. The Web services security policy is defined in the deployment descriptor extension (`ibm-webservicesclient-ext.xmi`) and the bindings are stored in the IBM binding extension (`ibm-webservicesclient-bnd.xmi`). However, the `${install_dir}/properties/ws-security.xml` file defines the default binding value for the client. If the binding information is not specified in the binding file, the run time reads the binding information from the default `${install_dir}/properties/ws-security.xml` file.

### ***Login mappings collection:***

Use this page to view a list of configurations for validating security tokens within incoming messages. Login mappings map an authentication method to a Java Authentication and Authorization Service (JAAS) login configuration to validate the security token. Four authentication methods are predefined in the WebSphere Application Server: BasicAuth, Signature, IDAssertion, and Lightweight Third Party Authentication (LTPA).

To view this administrative console page, complete the following steps:

1. Click **Server > Application Servers > *server\_name***.
2. Under Additional Properties, click **Web Services: Default bindings for Web Services Security > Login Mappings**.
3. Click **New** to create a login mapping or click **Delete** to delete a login mapping.

If you click **Update runtime**, the Web services security run time is updated with the default binding information, which is contained in the `ws-security.xml` file that was previously saved. After you specify the authentication method, the Java Authentication and Authorization Service (JAAS) configuration name, and the Callback Handler Factory class name on this panel, you must complete the following steps:

Click **Save** at the top of the administrative console. When you click **Save**, you return to the administrative console home panel.

Return to the Login Mappings collection panel and click **Update runtime**. When you click **Update runtime**, the configuration changes made to the other Web services also are updated in the Web services security run time.

#### *Authentication Method:*

Specifies the authentication method used for validating the security tokens.

The following authentication methods are available:

#### **BasicAuth**

The basic authentication method includes both a user name and a password in the security token. The information in the token is authenticated by the receiving server and used to create a credential.

#### **Signature**

The signature authentication method sends an X.509 certificate as a security token. For Lightweight Directory Access Protocol (LDAP) registries, the distinguished name (DN) is mapped

to a credential, which is based on the LDAP certificate filter settings. For local OS registries, the first attribute of the certificate, usually the common name (CN) is mapped directly to a user ID in the registry.

#### **IDAssertion**

The identity assertion method maps a trusted identity (ID) to a WebSphere Application Server credential. This authentication method only includes a user name in the security token. An additional token is included in the message for trust purposes. When the additional token is trusted, the IDAssertion token user name is mapped to a credential.

**LTPA** Lightweight Third Party Authentication (LTPA) validates an LTPA token.

*JAAS Configuration Name:*

Specifies the name of the Java Authentication and Authorization Service (JAAS) configuration.

*Callback Handler Factory Class Name:*

Specifies the name of the factory for the CallbackHandler class.

#### **Login mapping configuration settings:**

Use this page to specify the Java Authentication and Authorization Service (JAAS) login configuration settings used to validate security tokens within incoming messages.

To view this administrative console page, complete the following steps:

1. Click **Servers > Application Servers > server\_name**.
2. Under Additional Properties, click **Web Services: Default bindings for Web Services Security > Login Mappings > New**.

*Authentication Method:*

Specifies the method of authentication.

You can use any string, but the string must match the element in the service-level configuration. The following words are reserved and have special meanings:

#### **BasicAuth**

Uses both a user name and a password.

#### **IDAssertion**

Uses only a user name, but requires that additional trust is established on the receiving server using a TrustedIDEvaluator mechanism.

#### **Signature**

Uses the distinguished name (DN) of the signer.

**LTPA** Validates a token.

*JAAS Configuration Name:*

Specifies the name of the Java Authentication and Authorization Service (JAAS) configuration.

Specify your JAAS configurations using the administrative console by clicking **Security > JAAS Configuration > Application**.

*Callback Handler Factory Class Name:*

Specifies the name of the factory for the CallbackHandler class.



You must implement the `com.ibm.wsspi.wssecurity.auth.callback.CallbackHandlerFactory` class in this field.

**Default:** `com.ibm.wsspi.wssecurity.auth.callback.CallbackHandlerFactory`

*Token Type URI:*

Specifies the namespace Uniform Resource Identifiers (URI), which denotes the type of security token accepted.

If binary security tokens are accepted, the value denotes the `ValueType` attribute in the element. The `ValueType` element identifies the type of security token and its namespace. If Extensible Markup Language (XML) tokens are accepted, the value denotes the top-level element name of the XML token.

If the reserved words are specified previously in the Authentication Method field, this field is ignored.

**Data type:** Unicode characters except for non-ASCII characters, but including the number sign (#), the percent sign (%), and the square brackets ([ ]).

*Token Type Local Name:*

Specifies the local name of the security token type, for example, X509v3.

If binary security tokens are accepted, the value denotes the `ValueType` attribute in the element. The `ValueType` attribute identifies the type of security token and its namespace. If Extensible Markup Language (XML) tokens are accepted, the value denotes the top-level element name of the XML token.

If the reserved words are specified previously in the Authentication Method field, this field is ignored.

*Nonce Maximum Age:*

Specifies the time, in seconds, before the nonce time stamp expires. Nonce is a randomly generated value.

You must specify a minimum of 300 seconds for the Nonce Maximum Age field. However, the maximum value cannot exceed the number of seconds specified in the Nonce Cache Timeout field for either the server level or the cell level. You can specify the **Nonce Maximum Age** value for the server level:

1. Click **Servers > Application Servers > *server\_name***.
2. Under Additional Properties, click **Web Services: Default bindings for Web Services Security**.
3. Specify the **Nonce Maximum Age** value for the cell level by clicking **Security > Web Services > Properties**.

**Important:** The Nonce Maximum Age field on this panel is optional and only valid if the BasicAuth authentication method is specified. If you specify another authentication method and attempt to specify values for this field, the following error message displays and you must remove the specified value: Nonce is not supported for authentication methods other than BasicAuth.

If you specify BasicAuth, but do not specify values for the Nonce Maximum Age field, the Web services security run time searches for a Nonce Maximum Age value on the server level. If a value is not found on the server level, the run time searches the cell level. If a value is not found on either the server level or the cell level, the default is 300 seconds.

**Default** 300 seconds

**Range**

300 to Nonce Cache Timeout seconds

*Nonce Clock Skew:*

Specifies the clock skew value, in seconds, to consider when WebSphere Application Server checks the freshness of the message. Nonce is a randomly generated value.

You must specify a minimum of 0 seconds for the Nonce Clock Skew field. However, the maximum value cannot exceed the number of seconds specified in the Nonce Maximum Age on this Login Mappings panel.

**Important:** The Nonce Clock Skew field on this panel is optional and only valid if the BasicAuth authentication method is specified. If you specify another authentication method and attempt to specify values for this field, the following error message displays and you must remove the specified value: Nonce is not supported for authentication methods other than BasicAuth.

**Note:** If you specify BasicAuth, but do not specify values for the Nonce Clock Skew field, the Web services security run time searches for a Nonce Clock Skew value on the server level. If a value is not found on the server level, the run time searches the cell level. If a value is not found on either the server level or the cell level, the default is 0 seconds.

**Default**

0 seconds

**Range**

0 to Nonce Maximum Age seconds

## Configuring the client for request signing: Digitally signing message parts

Prior to completing these steps, read either of the following topics to become familiar with the Security Extensions tab and the Port Binding tab in the Web Services Client Editor within the Assembly Toolkit.

- Configuring the client security bindings using the WebSphere Application Server Toolkit
- Configuring the security bindings on a server acting as a client using the administrative console

These two tabs are used to configure the Web services security extensions and the Web services security bindings, respectively.

Complete the following steps to specify which message parts to digitally sign when configuring the client for request signing:

1. Launch the WebSphere Application Server Toolkit and click **Windows > Open Perspective > J2EE**.
2. Select the Web services-enabled Enterprise JavaBeans (EJB) or Web module.
3. In the Package Explorer window, locate the META-INF directory for an EJB module or the WEB-INF directory for a Web module.
4. Right-click the **webservicesclient.xml** file and click **Open With > Web Services Client Editor**.
5. Click the **Security Extensions** tab, which is located at the bottom of the Web Services Client Editor within the Assembly Toolkit.
6. Expand **Request Sender Configuration > Integrity**. Integrity refers to digital signature while confidentiality refers to encryption. Integrity decreases the risk of data modification while the data is transmitted across the Internet. For more information on digitally signing SOAP messages, see XML digital signature.
7. Indicate which parts of the message to sign by clicking **Add** and selecting **body**, **timestamp**, or **SecurityToken**. The following list contains descriptions of the message parts

**Body** The body is the user data portion of the message.

## Timestamp

The time stamp determines if the message is valid based on the time that the message is sent and then received. If **timestamp** is selected, proceed to the next step and select **Add Created Time Stamp** to add a time stamp to a message.

## Securitytoken

The security token authenticates the client. If this option is selected, the message is signed.

You can choose to digitally sign the message using a time stamp if **Add Created Time Stamp** is selected and configured. You can digitally sign the message using a security token if a login configuration authentication method is selected.

- Optional: Expand the **Add Created Time Stamp** section and select this option if you want a time stamp added to the message. You can specify an expiration time for the time stamp, which helps defend against replay attacks. The lexical representation for duration is the [ISO 8601] extended format *PnYnMnDTnHnMnS*, where:

- *nY* represents the number of years
- *nM* represents the number of months
- *nD* represents the number of days
- *T* is the date and time separator
- *nH* represents the number of hours
- *nM* represents the number of minutes
- *nS* represents the number of seconds. The number of seconds can include decimal digits to arbitrary precision.

For example, to indicate a duration of 1 year, 2 months, 3 days, 10 hours, and 30 minutes, the format is: *P1Y2M3DT10H30M*. Typically, you configure a message time stamp for about 10 to 30 minutes, for example, 10 minutes is represented as: *P0Y0M0DT0H10M0S*. The *P* character precedes time and date values.

**Important:** If you configure the client and server signing information correctly, but receive a Soap body not signed error when executing the client, you might need to configure the actor. You can configure the actor in the following locations on the client in the Web Services Client Editor within the Assembly Toolkit:

- Click **Security Extensions > Client Service Configuration Details** and indicate the actor information in the **Actor URI** field.
- Click **Security Extensions > Request Sender Configuration > Details** and indicate the actor information in the **Actor** field.

You must configure the same actor strings for the Web service on the server, which processes the request and sends the response back. Configure the actor in the following locations in the Web Services Editor within the WebSphere Application Server Toolkit:

- Click **Security Extensions > Server Service Configuration**.
- Click **Security Extensions > Response Sender Service Configuration Details > Details** and indicate the actor information in the **Actor** field.

The actor information on both the client and server must refer to the same exact string. When the actor fields on the client and server match, the request or response is acted upon instead of being forwarded downstream. The **actor** fields might be different when you have Web services acting as a gateway to other Web services. However, in all other cases, make sure that the actor information matches on the client and server. When Web services are acting as a gateway and they do not have the same actor configured as the request passing through the gateway, Web services do not process the message from a client. Instead, these Web services send the request downstream. The downstream process that contains the correct actor string processes the request. The same situation occurs for the response. Therefore, it is important that you verify that the appropriate client and server actor fields are synchronized.

After you specify which message parts to digitally sign when the client sends a message to a server.

Once you have specified which message parts to digitally sign, you must specify which method is used to digitally sign the message. See *Configuring the client for request signing: Choosing the digital signature method* for more information.

## Configuring the client for request signing: Choosing the digital signature method

Prior to completing these steps, read either of the following topics to become familiar with the Security Extensions tab and the Port Binding tab in the Web Services Client Editor within the Assembly Toolkit:

- *Configuring the client security bindings using the WebSphere Application Server Toolkit*
- *Configuring the security bindings on a server acting as a client using the administrative console*

These two tabs are used to configure the Web services security extensions and the Web services security bindings, respectively. You must specify which parts of the message sent by the client must be digitally signed. See *Configuring the client for request signing: Digitally signing message parts* for more information.

Complete the following steps to specify which message parts to digitally sign when configuring the client for request signing:

1. Launch the Assembly Toolkit and click **Windows > Open Perspective > J2EE**.
2. Select the Web services enabled EJB or Web module.
3. In the Package Explorer window, locate the META-INF directory for an EJB module or the WEB-INF directory for a Web module.
4. Right-click the **webservicesclient.xml** file and click **Open With > Web Services Client Editor**.
5. Click the **Port Binding** tab.
6. Expand **Security Request Sender Binding Configuration > Signing Information**.
7. Select **Edit** to view the signing information and select a digital signature method from the **Signature method algorithm** field. The following table describes the purpose of this information. Some of these definitions are based on the XML-Signature specification, which is located at the following Web site <http://www.w3.org/TR/xmldsig-core>.

Name	Purpose
<b>Canonicalization method algorithm</b>	Canonicalizes the <SignedInfo> element before the information is digested as part of the signature operation.
<b>Digest method algorithm</b>	Applies to the data after transforms are applied, if specified, to yield the <DigestValue> element. Signing the <DigestValue> element binds the resource content to the signer key. The algorithm selected for the client request sender configuration must match the algorithm selected in the server request receiver configuration.
<b>Signature method algorithm</b>	Converts the canonicalized <SignedInfo> element into the <SignatureValue> element. The algorithm selected for the client request sender configuration must match the algorithm selected in the server request receiver configuration.
<b>Signing key name</b>	Represents the key entry associated with the signing key locator. The key entry refers to an alias of the key, which is found in the key store and is used to sign the request.
<b>Signing key locator</b>	Represents a reference to a key locator implementation class that locates the correct key store where the alias and the certificate exist.

**Important:** If you configure the client and server signing information correctly, but receive a Soap body not signed error when executing the client, you might need to configure the actor. You can configure the actor in the following locations on the client in the Web Services Client Editor within the WebSphere Application Server Toolkit:

- Click **Security Extensions > Client Service Configuration Details** and indicate the actor information in the **Actor URI** field.
- Click **Security Extensions > Request Sender Configuration > Details** and indicate the actor information in the **Actor** field.

You must configure the same actor strings for the Web service on the server, which processes the request and sends the response back. Configure the actor in the following locations in the Web Services Editor within the WebSphere Application Server Toolkit:

- Click **Security Extensions > Server Service Configuration**.
- Click **Security Extensions > Response Sender Service Configuration Details > Details** and indicate the actor information in the **Actor** field.

The actor information on both the client and server must refer to the same exact string. When the actor fields on the client and server match, the request or response is acted upon instead of being forwarded downstream. The **actor** fields might be different when you have Web services acting as a gateway to other Web services. However, in all other cases, make sure that the actor information matches on the client and server. When Web services are acting as a gateway and they do not have the same actor configured as the request passing through the gateway, Web services do not process the message from a client. Instead, these Web services send the request downstream. The downstream process that contains the correct actor string processes the request. The same situation occurs for the response. Therefore, it is important that you verify that the appropriate client and server actor fields are synchronized.

You have specified which method is used to digitally sign a message when the client sends a message to a server.

After you configure the client to digitally sign the message, you must configure the server to verify the digital signature. See [Configuring the server for request digital signature verification: Verifying the message parts](#) for more information.

## **Configuring the server for request digital signature verification: verifying the message parts**

Prior to completing these steps, read either of the following topics to become familiar with the Security Extensions tab and the Binding Configurations tab in the Web services editor within the WebSphere Application Server Toolkit:

- [Configuring the server security bindings using the WebSphere Application Server Toolkit](#)
- [Configuring the server security bindings using the administrative console](#)

You can use these two tabs to configure the Web services security extensions and the Web services security bindings, respectively. Also, you must specify which parts of the message sent by the client must be digitally signed. See [Configuring the client for request signing: Digitally signing message parts](#) to determine which message parts are digitally signed. The message parts specified for the client request sender must match the message parts specified for the server request receiver.

Complete the following steps to configure the server for request digital signature verification. The steps describe how to modify the extensions to indicate which parts of the request to verify.

1. Launch the WebSphere Application Server Toolkit and either click **Windows > Open Prospective > Java** or **Windows > Open Prospective > Resource**.
2. Select the Web services-enabled Enterprise JavaBean (EJB) or Web module.

3. In the Package Explorer window, locate the META-INF directory for an EJB module or the WEB-INF directory for a Web module.
4. Right-click the `webservices.xml` file, and click **Open With > Web Services Editor**.
5. Click the **Security Extensions** tab in the Web services editor.
6. Expand the **Request Receiver Service Configuration Details > Required Integrity** section. Required integrity refers to the parts of the message that require digital signature verification. The purpose of digital signature verification is to make sure that the message parts have not been modified while transmitting across the Internet.
7. Indicate parts of the message to verify by clicking **Add**, and selecting one of the following three parts: body, timestamp, or SecurityToken. You can determine which parts of the message to verify by looking at the Web service request sender configuration in the client application. To view the Web service request sender configuration information in the Web services client editor, click the **Security Extensions** tab and expand **Request Sender Configuration > Integrity**. The following includes a list and description of the message parts:

**Body** This is the user data portion of the message.

**Timestamp**

The timestamp determines if the message is valid based on the time that the message is sent and then received. If timestamp option is selected, proceed to the next step to Add Created Time Stamp to the message.

**Securitytoken**

The security token authenticates the client. If the SecurityToken is selected, the message is signed.

8. Optional: Expand the **Add Received Time Stamp** section. The Add Received Time Stamp value indicates to validate the Add Created Time Stamp option configured by the client. You must select this option if you selected the Add Created Time Stamp on the client. The time stamp ensures message integrity by indicating the timeliness of the request. This option helps defend against replay attacks.

**Important:** If you configure the client and server signing information correctly, but receive a Soap body not signed error when executing the client, you might need to configure the actor. You can configure the actor in the following locations on the client in the Web Services Client Editor within the WebSphere Application Server Toolkit:

- Click **Security Extensions > Client Service Configuration Details** and indicate the actor information in the **Actor URI** field.
- Click **Security Extensions > Request Sender Configuration > Details** and indicate the actor information in the **Actor** field.

You must configure the same actor strings for the Web service on the server, which processes the request and sends the response back. Configure the actor in the following locations in the Web Services Editor within the WebSphere Application Server Toolkit:

- Click **Security Extensions > Server Service Configuration**.
- Click **Security Extensions > Response Sender Service Configuration Details > Details** and indicate the actor information in the **Actor** field.

The actor information on both the client and server must refer to the same exact string. When the actor fields on the client and server match, the request or response is acted upon instead of being forwarded downstream. The **actor** fields might be different when you have Web services acting as a gateway to other Web services. However, in all other cases, make sure that the actor information matches on the client and server. When Web services are acting as a gateway and they do not have the same actor configured as the request passing through the gateway, Web services do not process the message from a client. Instead, these Web services send the request downstream. The downstream process that contains the correct



actor string processes the request. The same situation occurs for the response. Therefore, it is important that you verify that the appropriate client and server actor fields are synchronized.

You have specified which message parts are digitally signed and must be verified by the server when the client sends a message to a server.

After you specify which message parts contain a digital signature that must be verified by the server, you must configure the server to recognize the digital signature method used to digitally sign the message. See *Configuring the server for request digital signature verification: Choosing the verification method* for more information.

## Configuring the server for request digital signature verification: choosing the verification method

Prior to completing these steps, read either of the following topics to become familiar with the Security Extensions tab and the Binding Configurations tab in the Web Services Editor within the WebSphere Application Server Toolkit:

- *Configuring the server security bindings using the WebSphere Application Server Toolkit*
- *Configuring the server security bindings using the administrative console*

You can use these two tabs to configure the Web services security extensions and Web services security bindings, respectively. You must specify which message parts contain digital signature information that must be verified by the server. See *Configuring the server for request digital signature verification: verifying the message parts*. The message parts specified for the client request sender must match the message parts specified for the server request receiver. Likewise, the digital signature method chosen for the client must match the digital signature method used by the server.

Complete the following steps to configure the server for request digital signature verification. The steps describe how to modify the extensions to indicate which digital signature method the server will use during verification.

1. Launch the WebSphere Application Server Toolkit and either click **Windows > Open Prospective > Java** or **Windows > Open Prospective > Resource**.
2. Select the Web services-enabled Enterprise JavaBean (EJB) or Web module.
3. In the Package Explorer window, locate the META-INF directory for an EJB module or the WEB-INF directory for a Web module.
4. Right-click the webservices.xml file, and click **Open With > Web Services Editor**.
5. Click the **Binding Configurations** tab.
6. Expand the **Security Request Receiver Binding Configuration Details > Signing Information** section.
7. Click **Edit** to edit the signing information. The signing information dialog is displayed, select or enter the following information:
  - Canonicalization method algorithm
  - Digest method algorithm
  - Signature method algorithm
  - Use certificate path reference
  - Trust anchor reference
  - Certificate store reference
  - Trust any certificate

For more conceptual information on digitally signing Simple Object Access Protocol (SOAP) messages, see *XML digital signature*. The following table describes the purpose for each of these selections. Some of the following definitions are based on the XML-Signature specification, which is located at the following Web address: <http://www.w3.org/TR/xmlsig-core>.



Name	Purpose
<b>Canonicalization method algorithm</b>	Canonicalizes the <SignedInfo> element before it is digested as part of the signature operation. The algorithm selected for the server request receiver configuration must match the algorithm selected in the client request sender configuration.
<b>Digest method algorithm</b>	Applies to the data after transforms are applied, if specified, to yield the <DigestValue> element. The signing of the <DigestValue> element binds resource content to the signer key. The algorithm selected for the server request receiver configuration must match the algorithm selected in the client request sender configuration.
<b>Signature method algorithm</b>	Converts the canonicalized <SignedInfo> element into the <SignatureValue> element. The algorithm selected for the server request receiver configuration must match the algorithm selected in the client request sender configuration.
<b>Use certificate path reference or Trust any certificate</b>	Validates a certificate or signature sent with a message. When a message is signed, the public key used to sign it is sent with the message. This public key or certificate might not be validated at the receiving end. By selecting <b>User certificate path reference</b> , you must configure a trust anchor reference and a certificate store reference to validate the certificate sent with the message. By selecting <b>Trust any certificate</b> , the signature is validated by the certificate sent with the message without the certificate itself being validated.
<b>Use certificate path reference: Trust anchor reference</b>	Refers to a key store that contains trusted, self-signed certificates and certificate authority (CA) certificates. These certificates are trusted certificates that you can use with any applications in your deployment.
<b>Use certificate path reference: Certificate store reference</b>	Contains a collection of X.509 certificates. These certificates are not trusted for all applications in your deployment, but might be used as an intermediary to validate certificates for an application. See

**Important:** If you configure the client and server signing information correctly, but receive a Soap body not signed error when executing the client, you might need to configure the actor. You can configure the actor in the following locations on the client in the Web Services Client Editor within the WebSphere Application Server Toolkit:

- Click **Security Extensions > Client Service Configuration Details** and indicate the actor information in the **Actor URI** field.
- Click **Security Extensions > Request Sender Configuration > Details** and indicate the actor information in the **Actor** field.

You must configure the same actor strings for the Web service on the server, which processes the request and sends the response back. Configure the actor in the following locations in the Web Services Editor within the WebSphere Application Server Toolkit:

- Click **Security Extensions > Server Service Configuration**.
- Click **Security Extensions > Response Sender Service Configuration Details > Details** and indicate the actor information in the **Actor** field.

The actor information on both the client and server must refer to the same exact string. When the actor fields on the client and server match, the request or response is acted upon instead

of being forwarded downstream. The **actor** fields might be different when you have Web services acting as a gateway to other Web services. However, in all other cases, make sure that the actor information matches on the client and server. When Web services are acting as a gateway and they do not have the same actor configured as the request passing through the gateway, Web services do not process the message from a client. Instead, these Web services send the request downstream. The downstream process that contains the correct actor string processes the request. The same situation occurs for the response. Therefore, it is important that you verify that the appropriate client and server actor fields are synchronized.

You have specified the method that the server uses to verify the digital signature in the message parts.

After you configure the client for request signing and the server for request digital signature verification, you must configure the server and the client to handle the response. Next, specify the response signing for the server. See *Configuring the server for response signing: Digitally signing message parts* for more information.

### Configuring the server for response signing: Digitally signing message parts

Prior to completing these steps, read either of the following topics to become familiar with the Security Extensions tab and the Binding Configurations tab in the Web Services Editor within the Assembly Toolkit:

- *Configuring the server security bindings using the WebSphere Application Server Toolkit*
- *Configuring the server security bindings using the administrative console*

These two tabs are used to configure the Web services security extensions and the Web services security bindings, respectively.

Complete the following steps to specify which message parts to digitally sign when configuring the server for response signing:

1. Launch the WebSphere Application Server Toolkit and click **Windows > Open Perspective > J2EE**.
2. Select the Web services-enabled Enterprise JavaBean (EJB) or Web module.
3. In the Package Explorer window, locate the META-INF directory for an EJB module or the WEB-INF directory for a Web module.
4. Right-click the **webservices.xml** file and click **Open With > Web Services Editor**.
5. Click the **Security Extensions** tab, which is located at the bottom of the Web Services Editor within the Assembly Toolkit.
6. Expand **Response Sender Service Configuration Details > Integrity**. Integrity refers to digital signature while confidentiality refers to encryption. Integrity decreases the risk of data modification while the data is transmitted across the Internet. For more information on digitally signing SOAP messages, see XML digital signature.
7. Indicate the parts of the message to sign by clicking **Add**, and selecting **Body**, **Timestamp**, or **Securitytoken**.

The following list contains descriptions of the message parts:

**Body** The body is the user data portion of the message.

#### Timestamp

The time stamp determines if the message is valid based on the time that the message is sent and then received. If **timestamp** is selected, proceed to the next step and click **Add Created Time Stamp**, which indicates that the time stamp is added to the message.

#### Securitytoken

The security token is used for authentication. If this option is selected, the authentication information is added to the message.

8. Optional: Expand the **Add Created Time Stamp** section. Select this option if you want a time stamp added to the message. You can specify an expiration time for the time stamp, which helps defend against replay attacks. The lexical representation for duration is the ISO 8601 extended format, *PnYnMnDTnHnMnS*, where:

- *nY* represents the number of years.
- *nM* represents the number of months.
- *nD* represents the number of days.
- *T* is the date and time separator.
- *nH* represents the number of hours.
- *nM* represents the number of minutes.
- *nS* represents the number of seconds. The number of seconds can include decimal digits to arbitrary precision.

For example, to indicate a duration of 1 year, 2 months, 3 days, 10 hours, and 30 minutes, the format is: P1Y2M3DT10H30M. Typically, you configure a message time stamp for about 10 to 30 minutes. 10 minutes is represented as: P0Y0M0DT0H10M0S. The *P* character precedes time and date values.

**Important:** If you configure the client and server signing information correctly, but receive a Soap body not signed error when executing the client, you might need to configure the actor. You can configure the actor in the following locations on the client in the Web Services Client Editor within the Assembly Toolkit:

- Click **Security Extensions > Client Service Configuration Details** and indicate the actor information in the **Actor URI** field.
- Click **Security Extensions > Request Sender Configuration > Details** and indicate the actor information in the **Actor** field.

You must configure the same actor strings for the Web service on the server, which processes the request and sends the response back. Configure the actor in the following locations in the Web Services Editor within the WebSphere Application Server Toolkit:

- Click **Security Extensions > Server Service Configuration**.
- Click **Security Extensions > Response Sender Service Configuration Details > Details** and indicate the actor information in the **Actor** field.

The actor information on both the client and server must refer to the same exact string. When the actor fields on the client and server match, the request or response is acted upon instead of being forwarded downstream. The **actor** fields might be different when you have Web services acting as a gateway to other Web services. However, in all other cases, make sure that the actor information matches on the client and server. When Web services are acting as a gateway and they do not have the same actor configured as the request passing through the gateway, Web services do not process the message from a client. Instead, these Web services send the request downstream. The downstream process that contains the correct actor string processes the request. The same situation occurs for the response. Therefore, it is important that you verify that the appropriate client and server actor fields are synchronized.

You have specified which message parts to digitally sign when the server sends a response to the client.

After you specifying which message parts to digitally sign, you must specify which method is used to digitally sign the message. See *Configuring the server for response signing: Choosing the digital signature method* for more information.

## Configuring the server for response signing: choosing the digital signature method

Prior to completing these steps, read either of the following topics to become familiar with the Security Extensions tab and the Binding Configurations tab in the Web Services Editor within the WebSphere Application Server Toolkit:

- Configuring the server security bindings using the WebSphere Application Server Toolkit
- Configuring the server security bindings using the administrative console

These two tabs are used to configure the Web services security extensions and the Web services security bindings, respectively.

Complete the following steps to specify which digital signature method to use when configuring the server for response signing:

1. Launch the Assembly Toolkit and click **Windows > Open Perspective > J2EE**.
2. Select the Web services-enabled Enterprise JavaBean (EJB) or Web module.
3. In the Package Explorer window, locate the META-INF directory for an EJB module or the WEB-INF directory for a Web module.
4. Right-click the **webservices.xml** file and click **Open With > Web Services Editor**.
5. Click the **Binding Configurations** tab.
6. Expand **Response Sender Binding Configuration Details > Signing Information**.
7. Click **Edit** to choose a signing method. The signing info dialog is displayed and either select or enter the following information:
  - **Canonicalization method algorithm**
  - **Digest method algorithm**
  - **Signature method algorithm**
  - **Signing key name**
  - **Signing key locator**

The following table describes the purpose of this information. Some of these definitions are based on the XML-Signature specification, which is located at the following address:  
<http://www.w3.org/TR/xmlsig-core>.

Name	Purpose
<b>Canonicalization method algorithm</b>	Canonicalizes the <SignedInfo> element before the information is digested as part of the signature operation. Use the same algorithm on the client response receiver. The algorithm selected for the server response sender configuration must match the algorithm selected in the client response receiver configuration.
<b>Digest method algorithm</b>	Applies to the data after transforms are applied, if specified, to yield the <DigestValue> element. Signing the <DigestValue> binds resource content to the signer key. The algorithm selected for the server response sender configuration must match the algorithm selected in the client response receiver configuration.
<b>Signature method algorithm</b>	Converts the canonicalized <SignedInfo> element into the <SignatureValue> element. The algorithm selected for the server response sender configuration must match the algorithm selected in the client response receiver configuration.

Name	Purpose
<b>Signing key name</b>	Represents the key entry associated with the signing key locator. The key entry refers to an alias of the key, which is found in the key store and is used to sign the request.
<b>Signing key locator</b>	Represents a reference to a key locator implementation class that locates the correct key store where the alias and certificate exists. For more information on configuring key locators, see any of the following files: <ul style="list-style-type: none"> <li>• Configuring key locators using the WebSphere Application Server Toolkit</li> <li>• Configuring key locators using the administrative console</li> </ul>

You have specified which method is used to digitally sign a message when the server sends a message to a client.

After you configure the server to digitally sign the response message, you must configure the client to verify the digital signature contained in the response message. See *Configuring the client for response digital signature verification: Verifying the message parts* for more information.

### **Configuring the client for response digital signature verification: Verifying the message parts**

Prior to completing these steps, read either of the following topics to become familiar with the Security Extensions tab and the Port Binding tab in the Web Services Client Editor within the Assembly Toolkit:

- Configuring the client security bindings using the WebSphere Application Server Toolkit
- Configuring the security bindings on a server acting as a client using the administrative console

You can use these two tabs to configure the Web services security extensions and the Web services security bindings, respectively.

Complete the following steps to configure the client for response digital signature verification. The steps describe how to modify the extensions to indicate which parts of the response to verify.

1. Launch the Assembly Toolkit and click **Windows > Open Perspective > J2EE**.
2. Select the Web services-enabled Enterprise JavaBean (EJB) or Web module.
3. In the Package Explorer window, locate the META-INF directory for an EJB module or the WEB-INF directory for a Web module.
4. Right-click the **webservicesclient.xml** file and click **Open With > Web Services Client Editor**.
5. Click the **Security Extensions** tab.
6. Expand the **Response Receiver Configuration > Required Integrity** section. Required integrity refers to parts that require digital signature verification. Digital signature verification decreases the risk that the message parts have been modified while the message is transmitted across the Internet.
7. Indicate the parts of the message that must be verified. You can determine which parts of the message to verify by looking at the Web service response sender configuration. Click **Add** and select one of the following parts:

**Body** The body is the user data portion of the message.

#### **Timestamp**

The time stamp determines if the message is valid based on the time that the message is sent and then received. If the time stamp option is selected, proceed to the next step to add a received time stamp to the message.

## Securitytoken

The security token authenticates the client. If Securitytoken option is selected, the message is signed.

- Optional: Expand the **Add Received Time Stamp** section. Select **Add Received Time Stamp** to add the received time stamp to the message.

**Important:** If you configure the client and server signing information correctly, but receive a Soap body not signed error when executing the client, you might need to configure the actor. You can configure the actor in the following locations on the client in the Web Services Client Editor within the WebSphere Application Server Toolkit:

- Click **Security Extensions > Client Service Configuration Details** and indicate the actor information in the **Actor URI** field.
- Click **Security Extensions > Request Sender Configuration > Details** and indicate the actor information in the **Actor** field.

You must configure the same actor strings for the Web service on the server, which processes the request and sends the response back. Configure the actor in the following locations in the Web Services Editor within the WebSphere Application Server Toolkit:

- Click **Security Extensions > Server Service Configuration**.
- Click **Security Extensions > Response Sender Service Configuration Details > Details** and indicate the actor information in the **Actor** field.

The actor information on both the client and server must refer to the same exact string. When the actor fields on the client and server match, the request or response is acted upon instead of being forwarded downstream. The **actor** fields might be different when you have Web services acting as a gateway to other Web services. However, in all other cases, make sure that the actor information matches on the client and server. When Web services are acting as a gateway and they do not have the same actor configured as the request passing through the gateway, Web services do not process the message from a client. Instead, these Web services send the request downstream. The downstream process that contains the correct actor string processes the request. The same situation occurs for the response. Therefore, it is important that you verify that the appropriate client and server actor fields are synchronized.

You have specified which message parts are digitally signed and must be verified by the client when the server sends a response message to the client.

After you specify which message parts contain a digital signature that must be verified by the client, you must configure the client to recognize the digital signature method used to digitally sign the message. See [Configuring the client for response digital signature verification: Choosing the verification method](#) for more information.

## Configuring the client for response digital signature verification: choosing the verification method

Prior to completing these steps, read either of the following topics to become familiar with the **Security Extensions** tab and the **Port Binding** tab in the Web Services Editor within the Application Server Toolkit:

- Configuring the server security bindings using the WebSphere Application Server Toolkit
- Configuring the server security bindings using the administrative console

You can use these two tabs to configure the Web Services Security extensions and Web Services Security bindings, respectively. Also, you must specify which message parts contain digital signature information that must be verified by the client. See [Configuring the client for response digital signature verification: verifying the message parts to specify which message parts are digitally signed by the server and must be verified by the client](#). The message parts specified for the server response sender must match the



message parts specified for the client response receiver. Likewise, the digital signature method chosen for the server must match the digital signature method used by the client.

Complete the following steps to configure the client for response digital signature verification. The steps describe how to modify the extensions to indicate which digital signature method the client will use during verification.

1. Launch the WebSphere Application Server Toolkit and either click **Windows > Open Prospective > Java** or **Windows > Open Prospective > Resource**.
2. Select the Web services enabled EJB or Web module.
3. In the Package Explorer window, locate the **META-INF** directory for an EJB module or the **WEB-INF** directory for a Web module.
4. Right-click the `webservicesclient.xml` file, select **Open With > Web Services Client Editor**.
5. Click the **Port Binding** tab.
6. Expand the **Security Response Receiver Binding Configuration > Signing Information** section.
7. Click **Edit** to select a digital signature method. The signing info dialog displays and either select or enter the following information:
  - **Canonicalization method algorithm**
  - **Digest method algorithm**
  - **Signature method algorithm**
  - **Signing key name**
  - **Signing key locator**

For more conceptual information on digitally signing SOAP messages, see XML digital signature. The following table describes the purpose for each of these selections. Some of the following definitions are based on the XML-Signature specification, which can be found at: <http://www.w3.org/TR/xmlsig-core>.

Name	Purpose
<b>Canonicalization method algorithm</b>	The canonicalization method algorithm is used to canonicalize the <code>&lt;SignedInfo&gt;</code> element before it is digested as part of the signature operation.
<b>Digest method algorithm</b>	The digest method algorithm is the algorithm applied to the data after transforms are applied, if specified, to yield the <code>&lt;DigestValue&gt;</code> . The signing of the <code>&lt;DigestValue&gt;</code> binds resource content to the signer key. The algorithm selected for the client response receiver configuration must match the algorithm selected in the server response sender configuration.
<b>Signature method algorithm</b>	The signature method is the algorithm that is used to convert the canonicalized <code>&lt;SignedInfo&gt;</code> element into the <code>&lt;SignatureValue&gt;</code> element. The algorithm selected for the client response receiver configuration must match the algorithm selected in the server response sender configuration.
<b>Use certificate path reference or Trust any certificate</b>	When a message is signed, the public key used to sign it is transmitted with the message. To validate this public key at the receiving end, configure a certificate path reference. By selecting <b>User certificate path reference</b> , you must configure a trust anchor reference and certificate store reference to validate the certificate sent with the message. By selecting <b>trust any certificate</b> , the signature is validated by the certificate sent with the message without the certificate itself being validated.



Name	Purpose
<b>Use certificate path reference: Trust anchor reference</b>	A trust anchor is a configuration that refers to a keystore that contains trusted, self-signed certificates and certificate authority (CA) certificates. These certificates are trusted certificates that you can use with any applications in your deployment.
<b>Use certificate path reference: Certificate store reference</b>	A certificate store is a configuration that has a collection of X.509 certificates. These certificates are not trusted for all applications in your deployment, but might be used as an intermediary to validate certificates for an application.

**Important:** If you configure the client and server signing information correctly, but receive a Soap body not signed error when executing the client, you might need to configure the actor. You can configure the actor in the following locations on the client in the Web Services Client Editor within the WebSphere Application Server Toolkit:

- Click **Security Extensions > Client Service Configuration Details** and indicate the actor information in the **Actor URI** field.
- Click **Security Extensions > Request Sender Configuration > Details** and indicate the actor information in the **Actor** field.

You must configure the same actor strings for the Web service on the server, which processes the request and sends the response back. Configure the actor in the following locations in the Web Services Editor within the WebSphere Application Server Toolkit:

- Click **Security Extensions > Server Service Configuration**.
- Click **Security Extensions > Response Sender Service Configuration Details > Details** and indicate the actor information in the **Actor** field.

The actor information on both the client and server must refer to the same exact string. When the actor fields on the client and server match, the request or response is acted upon instead of being forwarded downstream. The **actor** fields might be different when you have Web services acting as a gateway to other Web services. However, in all other cases, make sure that the actor information matches on the client and server. When Web services are acting as a gateway and they do not have the same actor configured as the request passing through the gateway, Web services do not process the message from a client. Instead, these Web services send the request downstream. The downstream process that contains the correct actor string processes the request. The same situation occurs for the response. Therefore, it is important that you verify that the appropriate client and server actor fields are synchronized.

You have specified which method the client uses to verify the digital signature in the message parts.

After you configure the server for response signing and the client for request digital signature verification, verify that you have configured the client and the server to handle the message request.

## Configuring the client security bindings using the Assembly Toolkit

When configuring a client for Web services security, the bindings describe how to run the security specifications found in the extensions. Use the Web services client editor within the Assembly Toolkit to include the binding information in the client enterprise archive (EAR) file.

You can configure the client-side bindings from a pure client accessing a Web service or from a Web service accessing a downstream Web service. This document focuses on the pure client situation. However, the concepts, and in most cases the steps, also apply when a Web service is configured to communicate downstream to another Web service that has client bindings. Complete the following steps to edit the security bindings on a pure client (or server acting as a client) using the Assembly Toolkit:

1. Import the Web services client EAR file into the Assembly Toolkit. When you edit the client bindings on a server acting as a client, the same basic steps apply. Complete the following steps to import your client EAR file into the Assembly Toolkit. Refer to the Assembly Toolkit documentation for additional information.
  - a. Download and install the Assembly Toolkit. You can download the Assembly Toolkit from the following Web site: [http://www.ibm.com/support/docview.wss?rs=180&context=SSEQTP&q=ASTK&uid=swg24005125&loc=en\\_US&cs=utf-8&lang=en+en](http://www.ibm.com/support/docview.wss?rs=180&context=SSEQTP&q=ASTK&uid=swg24005125&loc=en_US&cs=utf-8&lang=en+en)
  - b. Start the Assembly Toolkit and open the Java perspective, by clicking **Window > Open Perspective > J2EE**.
  - c. Import the client EAR file by clicking **File > Import > EAR file**.
  - d. Click **Next**.
  - e. Enter the path name to the EAR file in the **EAR File** field or click **Browse** to locate the file.
  - f. Enter the project name in the **Project name** field.
  - g. Click **Finish**.
2. Open the Web services client editor within the Assembly Toolkit to begin editing the client bindings. To access the client bindings using the Assembly Toolkit, complete the following steps:
  - a. Open the Navigator by clicking **Window > Show View > Navigator**.
  - b. Expand your application Java archive (JAR) file from the Navigator.
  - c. Expand the J2EE client application (appClientModule, ejbModule, or WebContent), which are included in the client JAR package that you selected.
  - d. Expand the META-INF directory and locate the webservicessclient.xml file.
  - e. Right-click the webservicessclient.xml file and click **Open With > Web Services Client Editor**. In the Web services client editor (for webservicessclient.xml and outbound requests and inbound responses Web services configuration), there are several tabs at the bottom of the editor including References, Handlers, Security Extensions, Web Services Client Binding, and Port Binding. The security extensions are edited using the **Security Extensions** tab. The security bindings are edited using the Port Binding tab.
3. On the **Security Extensions** tab, select the port qualified name bindings that you want to configure. The Web services security extensions are configured for outbound requests and inbound responses. You need to configure the following information for Web services security extensions. These topics are discussed in more detail in other sections of the documentation.

#### **Request sender configuration details**

##### **Details**

Configuring the client for request signing: digitally signing message parts

##### **Integrity**

Configuring the client for request signing: digitally signing message parts

##### **Confidentiality**

Configuring the client for request encryption: encrypting the message parts

##### **Login Config**

###### **BasicAuth**

Configuring the client for basicauth authentication: specifying the method

###### **IDAssertion**

Configuring the client for identity assertion authentication: specifying the method

###### **Signature**

Configuring the client for signature authentication: specifying the method

###### **LTPA**

Configuring the client for LTPA token authentication: specifying LTPA token authentication

**ID Assertion**

Configuring the client for identity assertion authentication: specifying the method

**Add Created Time Stamp**

Configuring the client for request signing: digitally signing message parts

**Response receiver configuration details****Required Integrity**

Configuring the client for response digital signature verification: verifying the message parts

**Required Confidentiality**

Configuring the client for response decryption: decrypting message parts

**Add Received Time Stamp**

Configuring the client for response digital signature verification: verifying the message parts

4. From the **Port Binding** tab, select the port qualified name binding that you want to configure. The Web services security bindings are configured for outbound requests and inbound responses. You need to configure the following information for Web services security bindings. These topics are discussed in more details in other sections of the documentation.

**Security request sender binding configuration****Signing information**

Configuring the client for request signing: choosing the digital signature method

**Encryption information**

Configuring the client for request encryption: choosing the encryption method

**Key locators**

Configuring key locators using the Assembly Toolkit

**Login binding****Basic auth**

Configuring the client for basicauth authentication: collecting the authentication information

**ID assertion**

Configuring the client for identity assertion: Collecting the authentication method

**Signature**

Configuring the client for signature authentication: collecting the authentication information

**LTPA**

Configuring the client for LTPA token authentication: Collecting the authentication method information

**Security response receiver binding configuration****Signing information**

Configuring the client for response digital signature verification: choosing the verification method

**Encryption information**

Configuring the client for response decryption: choosing a decryption method

**Trust anchor**

Configuring trust anchors using the Assembly Toolkit

**Certificate store list**

Configuring the client-side collection certificate store using the Application Server Toolkit

**Key locators**

Configuring key locators using the Assembly Toolkit

**Important:** When configuring the security request sender binding configuration, you must synchronize the information used to perform the specified security with the security request receiver binding configuration, which is configured in the server EAR file. These two configurations must be synchronized in all respects because there is no negotiation during run time to determine the requirements of the server.

For example, when configuring the encryption information in the security request sender binding Configuration, you must use the public key from the server for encryption. Therefore, the key locator that you choose must contain the public key from the server configuration. The server must contain the private key to decrypt the message. This example illustrates the important relationship between the client and server configuration. Additionally, when configuring the security response receiver binding configuration, the server must send the response using security information known by this client security response receiver binding configuration.

The following table shows the related configurations between the client and the server. The client request sender and the server request receiver are relative configurations that must be synchronized with each other. The server response sender and the client response receiver are related configurations that must be synchronized with each other. Note that the related configurations are end points for any request or response. One end point must communicate its actions with the other end point because run time requirements are not negotiated.

Table 7. Related configurations

Client configuration	Server configuration
Request sender	Request receiver
Response receiver	Response sender

## Configuring the security bindings on a server acting as a client using the administrative console

When configuring a client for Web services security, the bindings describe how to run the security specifications found in the extensions. Use the Web Services Client Editor within the WebSphere Application Server Toolkit to include the binding information in the client enterprise archive (EAR) file.

You can configure the client-side bindings from a pure client accessing a Web service or from a Web service accessing a downstream Web service. Complete the following steps to find the location in which to edit the client bindings from a Web service that is running on the server. When a Web service communicates with another Web service, you must configure client bindings to access the downstream Web service.

1. Deploy the Web service using the WebSphere Application Server administrative console by clicking **Applications > Install New Application**. You can access the administrative console by typing `http://localhost:9090/admin` in your Web browser unless you have changed the port number. For more information on installing an application, see *Installing a new application*
2. Click **Applications > Enterprise Applications > *application\_name***.
3. Under Related Items, click either **Web Modules** or **EJB Modules**, depending upon which type of service is the client to the downstream service.
  - For Web modules, click the Web archive (WAR) file that you configured as the client.
  - For Enterprise JavaBeans (EJB) modules, click the Java archive (JAR) file that you configured as the client.
4. Click the name of the WAR or JAR file.
5. Under Additional Properties, click **Web Services: Client Security Bindings**. A table displays with the following columns:
  - Component Name

- Port
- Web Service
- Request Sender Binding
- Request Receiver Binding
- HTTP Basic Authentication
- HTTP SSL Configuration

For Web services security, you must edit the request sender binding and response receiver binding configurations. You can use the defaults for some of the information at the server level . Default bindings are convenient because you can configure commonly reused elements such as key locators once and then reference their aliases in the application bindings.

6. View the default bindings for the server using the administrative console by clicking **Servers > Application Server > server1**. Under Additional Properties, click **Web Services: Default bindings for Web Services Security**. You can configure the following sections. These topics are discussed in more detail in other sections of the documentation.

- Request sender binding
  - Signing information
  - Encryption information
  - Key locators
  - Login bindings
- Response receiver binding
  - Signing information
  - Encryption information
  - Trust anchors
  - Collection certificate store
  - Key locators

**Important:** When configuring the security request sender binding configuration, you must synchronize the information used to perform the specified security with the security request receiver binding configuration, which is configured in the server EAR file. These two configurations must be synchronized in all respects because there is no negotiation during run time to determine the requirements of the server. For example, when configuring the encryption information in the security request sender binding configuration, you must use the public key from the server for encryption. Therefore, the key locator that you choose must contain the public key from the server configuration. The server must contain the private key to decrypt the message. This example illustrates the important relationship between the client and server configuration. Additionally, when configuring the security response receiver binding configuration, the server must send the response using security information known by this client security response receiver binding configuration.

The following table shows the related configurations between the client and the server. The client request sender and the server request receiver are relative configurations that must be synchronized with each other. The server response sender and the client response receiver are related configurations that must be synchronized with each other. Note that related configurations are end points for any request or response. One end point must communicate its actions with the other end point because run time requirements are not required.

*Table 8. Related configurations*

Client configuration	Server configuration
Request sender	Request receiver
Response receiver	Response sender

## Configuring the server security bindings using the Assembly Toolkit

Create an Enterprise JavaBean (EJB) file Java archive (JAR) file or a Web archive (WAR) file containing the security binding file (`ibm-webservices-bnd.xmi`) and the security extension file (`ibm-webservices-ext.xmi`). If this archive is acting as a client to a downstream service, you also need the client-side binding file (`ibm-webservicesclient-bnd.xmi`) and the client-side extension file (`ibm-webservicesclient-ext.xmi`). These files are generated using the WSDL2Java command. You can edit these files using the Web Services Editor in the Assembly Toolkit.

When configuring server-side security for Web services security, the security extensions configuration specifies what security is performed, the security bindings configuration indicates how to perform what is specified in the security extensions configuration. You can use the defaults for some elements at the cell and server levels in the bindings configuration, including key locators, trust anchors, the collection certificate store, trusted ID evaluators, and login mappings and reference these elements from the WAR and JAR binding configurations.

Prior to importing the Web services enterprise archive (EAR) file into the Assembly Toolkit, make sure that you have already run the `wsdl2java` command on your Web service to enable your J2EE application. You must import the Web services enterprise archive (EAR) file into the Assembly Toolkit. Complete the following steps to import your EAR file into the Assembly Toolkit:

1. Download, install, and launch the Assembly Toolkit.
2. Open the J2EE perspective by clicking **Window > Open Perspective > Other > J2EE**.
3. Import the application EAR file by clicking **File > Import > EAR file**.
4. Click **Next** and indicate both the EAR file name in the **EAR File** field and the project name in the **Project name** field.
5. Click **Finish**.

Refer to Assembly Toolkit documentation for more information.

Open the Web services editor in the Assembly Toolkit to begin editing the server security extensions and bindings. The following steps can locate the server security extensions and bindings. Other tasks specify how to configure each section of the extensions and bindings in more detail.

1. Expand your application module from the Navigator. If the Navigator is not shown, you can open it by clicking **Window > Show View > Navigator**.
2. If your application is a Web archive (WAR) file, perform the following steps:
  - a. Expand the **WebContent > WEB-INF** section.
  - b. Locate the **webservices.xml** file. The `webservices.xml` file represents the server-side (inbound) Web services configuration. The `webservicesclient.xml` file represents the client-side (outbound) Web services configuration.
    - 1) Configure the server for inbound requests and outbound responses security configuration by right-clicking the **webservices.xml** file and clicking **Open With > Web Services Editor**.
    - 2) Configure the client for outbound requests and inbound responses security configuration by right-clicking the **webservicesclient.xml** file and clicking **Open With > Web Services Client Editor**. For more information, see *Configuring the client security bindings using the Assembly Toolkit*.
3. If your application is an EJB Application (JAR) file, perform the following steps:
  - a. Expand the **ejbModule > META-INF** section.
  - b. Locate the **webservices.xml** file. The `webservices.xml` file represents the server-side (inbound) Web services configuration. The `webservicesclient.xml` file represents the client-side (outbound) Web services configuration.
    - 1) Configure the server for inbound requests and outbound responses security configuration by right-clicking the **webservices.xml** file and clicking **Open With > Web Services Editor**.



- 2) Configure the client for outbound requests and inbound responses security configuration by right-clicking the **webservicesclient.xml** file and clicking **Open With > Web Services Client Editor**. For more information, see *Configuring the client security bindings using the Assembly Toolkit*.
4. In the Web services editor (for the `webservices.xml` file and inbound requests and outbound responses Web services configuration), there are several tabs at the bottom of the editor including Web Services, Port Components, Handlers, Security Extensions, Bindings, and Binding Configurations. The security extensions are edited using the Security Extensions tab. The security bindings are edited using the Security Bindings tab.
  - a. Click the **Security Extensions** tab and select the port component binding to edit. The Web services security extensions are configured for inbound requests and outbound responses. You need to configure the following information for Web services security extensions. These topics are discussed in more detail in other topics in the documentation.

#### **Request receiver service configuration details**

##### **Required integrity**

Configuring the server for request digital signature verification: verifying the message parts

##### **Required confidentiality**

Configuring the server for request decryption: decrypting message parts

##### **Login config**

###### **Basic auth**

Configuring the server to handle basicauth authentication

###### **ID assertion**

Configuring the server to handle identity assertion authentication

###### **Signature**

Configuring the server to handle signature authentication

**LTPA** Configuring the server to handle LTPA token authentication

##### **Add received time stamp**

Configuring the server for request digital signature verification: verifying the message parts

#### **Response sender service configuration details**

##### **Details**

Configuring the server for response signing: digitally signing message parts

##### **Integrity**

Configuring the server for response signing: digitally signing message parts

##### **Confidentiality**

Configuring the server for response encryption: encrypting message parts

##### **Add created time stamp-**

Configuring the server for response signing: digitally signing message parts

- b. Click the **Binding Configurations** tab and select the port component binding to edit. The Web services security bindings are configured for inbound requests and outbound responses. You need to configure the following information for Web services security bindings. These topics are discussed in more details in other topics in the documentation.

#### **Response receiver binding configuration details**

##### **Signing Information**

Configuring the server for request digital signature verification: choosing the verification method

##### **Encryption Information**

Configuring the server for request decryption: choosing the decryption method



**Trust Anchor**

Configuring trust anchors using the Assembly Toolkit

**Certificate Store List**

Configuring the server-side collection certificate store using the Assembly Toolkit

**Key Locators**

Configuring key locators using the Assembly Toolkit

**Login Mapping****Basic auth**

Configuring the server to validate basicauth authentication information

**ID assertion**

Configuring the server to validate identity assertion authentication information

**Signature**

Configuring the server to validate signature authentication information

**LTPA**

Configuring the server to validate LTPA token authentication information

**Trusted ID Evaluator****Trusted ID Evaluator Reference****Response sender binding configuration details****Signing information**

Configuring the server for response signing: choosing the digital signature method

**Encryption information**

Configuring the server for response encryption: choosing the encryption method

**Key Locators**

Configuring key locators using the Assembly Toolkit

## Configuring the server security bindings using the administrative console

Create an Enterprise JavaBean (EJB) file Java archive (JAR) file or Web archive (WAR) file containing the security binding file (`ibm-webservices-bnd.xmi`) and the security extension file (`ibm-webservices-ext.xmi`). If this archive is acting as a client to a downstream service, you also need the client-side binding file (`ibm-webservicesclient-bnd.xmi`) and the client-side extension file (`ibm-webservicesclient-ext.xmi`). These files are generated using the WSDL2Java command. You can edit these files using the Web Services Editor in the Assembly Toolkit.

When configuring server-side security for Web services security, the security extensions configuration specifies what security is to be performed while the security bindings configuration indicates how to perform what is specified in the security extensions configuration. You can use the defaults for some elements at the cell and server levels in the bindings configuration, including key locators, trust anchors, the collection certificate store, trusted ID evaluators, and login mappings and reference them from the WAR and JAR binding configurations.

The following steps describe how to edit bindings for a Web service after these bindings are deployed on a server. When one Web service communicates with another Web service, you also must configure the client bindings to access the downstream Web service.

1. Deploy the Web service using the WebSphere Application Server administrative console. The Administrative Console is accessible by typing `http://localhost:9090/admin` in a Web browser. After you log into the administration console, click **Applications > Install New Application** to deploy the Web service. For more information, see *Installing a new application*.
2. After you deploy the Web service, click **Applications > Enterprise Applications > *application\_name***.

3. Under Related Items, click either **Web Modules** or **EJB Modules** depending on which service you want to configure.
  - a. If you select **Web Modules**, click the WAR file that you want to edit.
  - b. If you select **EJB Modules**, click the JAR file that you want to edit.
4. After you select a WAR or JAR file, under **Additional Properties**, click **Web Services: Client Security Bindings** for outbound requests and inbound responses. Click **Web Services: Server Security Bindings** for inbound requests and outbound responses.
5. If you click **Web Services: Server Security Bindings**, the following sections can be configured. These topics are discussed in more detail in other sections of the documentation.
  - Request receiver binding
    - Signing Information
    - Encryption Information
    - Trust anchors
    - Collection certificate store
    - Key locators
    - Trusted ID evaluators
    - Login mappings
  - Response sender binding
    - Signing Information
    - Encryption Information
    - Key locators

## XML encryption

XML encryption is a specification developed by World Wide Web (WWW) Consortium (W3C) in 2002 that contains the steps to encrypt data, the steps to decrypt encrypted data the XML syntax to represent encrypted data, the information used to decrypt the data, and a list of encryption algorithms such as triple DES, AES, and RSA.

You can apply XML encryption to an XML element, XML element content, and arbitrary data, including an XML document. For example, suppose that you need to encrypt the CreditCard element shown in the example 1.

### Example 1: Sample XML document

```
<PaymentInfo xmlns='http://example.org/paymentv2'>
 <Name>John Smith</Name>
 <CreditCard Limit='5,000' Currency='USD'>
 <Number>4019 2445 0277 5567</Number>
 <Issuer>Example Bank</Issuer>
 <Expiration>04/02</Expiration>
 </CreditCard>
</PaymentInfo>
```

### Example 2: XML document with a common secret key

Example 2 shows the XML document after encryption. The EncryptedData element represents the encrypted CreditCard element. The EncryptionMethod element describes the applied encryption algorithm, which is triple DES in this example. The KeyInfo element contains the information to retrieve a decryption key, which is a KeyName element in this example. The CipherValue element contains the ciphertext obtained by serializing and encrypting the CreditCard element.

```
<PaymentInfo xmlns='http://example.org/paymentv2'>
 <Name>John Smith</Name>
 <EncryptedData Type='http://www.w3.org/2001/04/xmlenc#Element'
 xmlns='http://www.w3.org/2001/04/xmlenc#'>
```

```

 <EncryptionMethod
 Algorithm='http://www.w3.org/2001/04/xmlenc#tripledes-cbc' />
 <KeyInfo xmlns='http://www.w3.org/2000/09/xmldsig#'>
 <KeyName>John Smith</KeyName>
 </KeyInfo>
 <CipherData>
 <CipherValue>ydUNqHkMrD...</CipherValue>
 </CipherData>
 </EncryptedData>
</PaymentInfo>

```

### Example 3: XML document encrypted with the public key of the recipient

In example 2, it is assumed that both the sender and recipient have a common secret key. If the recipient has a public and private key pair, which is most likely the case, the CreditCard element can be encrypted as shown in example 3. The EncryptedData element is the same as the EncryptedData element found in Example 2. However, the KeyInfo element contains an EncryptedKey .

```

<PaymentInfo xmlns='http://example.org/paymentv2'>
 <Name>John Smith</Name>
 <EncryptedData Type='http://www.w3.org/2001/04/xmlenc#Element'
 xmlns='http://www.w3.org/2001/04/xmlenc#'>
 <EncryptionMethod
 Algorithm='http://www.w3.org/2001/04/xmlenc#tripledes-cbc' />
 <KeyInfo xmlns='http://www.w3.org/2000/09/xmldsig#'>
 <EncryptedKey xmlns='http://www.w3.org/2001/04/xmlenc#'>
 <EncryptionMethod
 Algorithm='http://www.w3.org/2001/04/xmlenc#rsa-1_5' />
 <KeyInfo xmlns='http://www.w3.org/2000/09/xmldsig#'>
 <KeyName>Sally Doe</KeyName>
 </KeyInfo>
 <CipherData>
 <CipherValue>yMTEy0TA1M...</CipherValue>
 </CipherData>
 </EncryptedKey>
 </KeyInfo>
 <CipherData>
 <CipherValue>ydUNqHkMrD...</CipherValue>
 </CipherData>
 </EncryptedData>
</PaymentInfo>

```

### XML Encryption in the WSS-Core

WSS-Core specification is under development by Organization for the Advancement of Structured Information Standards (OASIS). The specification describes enhancements to Simple Object Access Protocol (SOAP) messaging to provide quality of protection through message integrity, message confidentiality, and single message authentication. The message confidentiality is realized by encryption based on XML Encryption.

The WSS-Core specification supports encryption of any combination of body blocks, header blocks, their sub-structures, and attachments of a SOAP message. The specification also requires that when you encrypt parts of a SOAP message, you prepend a reference from the security header block to the encrypted parts of the message. The reference can be a clue for a recipient to identify which encrypted parts of the message to decrypt.

The XML syntax of the reference varies according to what information is encrypted and how it is encrypted. For example, suppose that the CreditCard element in example 4 is encrypted with either a common secret key or the public key of the recipient.

### Example 4: Sample SOAP message

```

<SOAP-ENV:Envelope
 SOAP-ENV:encodingStyle='http://schemas.xmlsoap.org/soap/encoding/'
 xmlns:SOAP-ENV='http://schemas.xmlsoap.org/soap/envelope/'>
 <SOAP-ENV:Body>
 <PaymentInfo xmlns='http://example.org/paymentv2'>
 <Name>John Smith</Name>
 <CreditCard Limit='5,000' Currency='USD'>
 <Number>4019 2445 0277 5567</Number>
 <Issuer>Example Bank</Issuer>
 <Expiration>04/02</Expiration>
 </CreditCard>
 </PaymentInfo>
 </SOAP-ENV:Body>
</SOAP-ENV:Envelope>

```

The resulting SOAP messages are shown in Examples 5 and 6. In these example, the ReferenceList and EncryptedKey elements are used as references, respectively.

### Example 5: SOAP message encrypted with a common secret key

```

<SOAP-ENV:Envelope
 SOAP-ENV:encodingStyle='http://schemas.xmlsoap.org/soap/encoding/'
 xmlns:SOAP-ENV='http://schemas.xmlsoap.org/soap/envelope/'>
 <SOAP-ENV:Header>
 <Security SOAP-ENV:mustUnderstand='1'
 xmlns='http://schemas.xmlsoap.org/ws/2003/06/secext'>
 <ReferenceList xmlns='http://www.w3.org/2001/04/xmlenc#'>
 <DataReference URI='#ed1' />
 </ReferenceList>
 </Security>
 </SOAP-ENV:Header>
 <SOAP-ENV:Body>
 <PaymentInfo xmlns='http://example.org/paymentv2'>
 <Name>John Smith</Name>
 <EncryptedData Id='ed1'
 Type='http://www.w3.org/2001/04/xmlenc#Element'
 xmlns='http://www.w3.org/2001/04/xmlenc#'>
 <EncryptionMethod
 Algorithm='http://www.w3.org/2001/04/xmlenc#tripleDES-cbc' />
 <KeyInfo xmlns='http://www.w3.org/2000/09/xmldsig#'>
 <KeyName>John Smith</KeyName>
 </KeyInfo>
 <CipherData>
 <CipherValue>ydUNqHkMrD...</CipherValue>
 </CipherData>
 </EncryptedData>
 </PaymentInfo>
 </SOAP-ENV:Body>
</SOAP-ENV:Envelope>

```

### Example 6: SOAP message encrypted with the public key of the recipient

```

<SOAP-ENV:Envelope
 SOAP-ENV:encodingStyle='http://schemas.xmlsoap.org/soap/encoding/'
 xmlns:SOAP-ENV='http://schemas.xmlsoap.org/soap/envelope/'>
 <SOAP-ENV:Header>
 <Security SOAP-ENV:mustUnderstand='1'
 xmlns='http://schemas.xmlsoap.org/ws/2003/06/secext'>
 <EncryptedKey xmlns='http://www.w3.org/2001/04/xmlenc#'>
 <EncryptionMethod
 Algorithm='http://www.w3.org/2001/04/xmlenc#rsa-1_5' />
 <KeyInfo xmlns='http://www.w3.org/2000/09/xmldsig#'>
 <KeyName>Sally Doe</KeyName>
 </KeyInfo>
 <CipherData>
 <CipherValue>yMTEy0TA1M...</CipherValue>
 </CipherData>
 </EncryptedKey>
 </Security>
 </SOAP-ENV:Header>
 <SOAP-ENV:Body>
 <PaymentInfo xmlns='http://example.org/paymentv2'>
 <Name>John Smith</Name>
 <EncryptedData Id='ed1'
 Type='http://www.w3.org/2001/04/xmlenc#Element'
 xmlns='http://www.w3.org/2001/04/xmlenc#'>
 <EncryptionMethod
 Algorithm='http://www.w3.org/2001/04/xmlenc#tripleDES-cbc' />
 <KeyInfo xmlns='http://www.w3.org/2000/09/xmldsig#'>
 <KeyName>John Smith</KeyName>
 </KeyInfo>
 <CipherData>
 <CipherValue>ydUNqHkMrD...</CipherValue>
 </CipherData>
 </EncryptedData>
 </PaymentInfo>
 </SOAP-ENV:Body>
</SOAP-ENV:Envelope>

```

```

 <ReferenceList>
 <DataReference URI='#ed1' />
 </ReferenceList>
 </EncryptedKey>
 </Security>
 </SOAP-ENV:Header>
 <SOAP-ENV:Body>
 <PaymentInfo xmlns='http://example.org/paymentv2'>
 <Name>John Smith</Name>
 <EncryptedData Id='ed1'
 Type='http://www.w3.org/2001/04/xmlenc#Element'
 xmlns='http://www.w3.org/2001/04/xmlenc#'>
 <EncryptionMethod
 Algorithm='http://www.w3.org/2001/04/xmlenc#tripleDES-cbc' />
 <CipherData>
 <CipherValue>ydUNqHkMrD...</CipherValue>
 </CipherData>
 </EncryptedData>
 </PaymentInfo>
 </SOAP-ENV:Body>
</SOAP-ENV:Envelope>

```

### Relationship to digital signature

The WSS-Core specification also provides message integrity, which is realized by a digital signature based on the XML-Signature specification.

A combination of encryption and digital signature over common data introduces cryptographic vulnerabilities.

## Securing Web services using XML encryption

WebSphere Application Server provides several different methods to secure your Web services. Extensible Markup Language (XML) encryption is one of these methods. You can secure your Web services using any of the following methods:

- XML digital signature
- XML encryption
- Basicauth authentication
- Identity assertion authentication
- Signature authentication
- Pluggable token

XML encryption enables you to encrypt an XML element, the content of an XML element, or arbitrary data such as an XML document. Like XML digital signature, a message is sent by the client as the request sender to the server as the request receiver. The response is sent by the server as the response sender to the client as the request receiver. Unlike XML digital signature, which verifies the authenticity of the sender, XML encryption scrambles the message content using a key, which can be unscrambled by a receiver that possesses the same key. You can use XML encryption in conjunction with XML digital signature to scramble the content while verifying the authenticity of the message sender.

To use XML encryption to secure Web services, you must use the WebSphere Application Server Toolkit, which is available at the following Web

site:[http://www.ibm.com/support/docview.wss?rs=180&context=SSEQTP&q=ASTK&uid=swg24005125&loc=en\\_US&cs=utf-8&lang=en+en](http://www.ibm.com/support/docview.wss?rs=180&context=SSEQTP&q=ASTK&uid=swg24005125&loc=en_US&cs=utf-8&lang=en+en)

Complete the following steps.

1. Specify the encryption settings for the request sender. The message parts and the encryption method settings chosen for the request sender on the client must match the message parts and the method settings chosen for the request receiver on the server. To specify the encryption settings for the request sender:
  - a. Configure the client for request encryption: Encrypting the message parts.
  - b. Configure the client for request encryption: Choosing the encryption method.
2. Specify the encryption settings for the request receiver. The decryption settings chosen for the request receiver must match the encryption settings chosen for the request sender.  
To specify the decryption settings for the request receiver:
  - a. Configure the server for request decryption: Decrypting message parts.
  - b. Configure the server for request decryption: Choosing the decryption method.
3. Specify the encryption settings for the response sender. The message parts and the encryption method settings chosen for the response sender on the server must match the message parts and the method settings chosen for the response receiver on the client. To specify the encryption settings for the response sender:
  - a. Configure the server for response encryption: Encrypting message parts.
  - b. Configure the server for response encryption: Choosing the encryption method.
4. Specify the encryption settings for the response receiver.

**Remember:** The decryption settings chosen for the response receiver must match the encryption settings chosen for the response sender.

To specify the decryption settings for the response receiver, complete the following steps:

- a. Configure the client for response decryption: decrypting message parts.
- b. Configure the client for response decryption: choosing the decryption method.

After completing these steps, you have secured your Web services using XML encryption.

## Encryption information collection

Use this page to specify the configuration for the encrypting and decrypting parameters. This configuration is used to encrypt and decrypt parts of the message, including the body and user name token.

To view this administrative console page, complete the following steps:

1. Click **Enterprise Applications** > *application\_name*.
2. Under Related Items, click **Web Module**.
3. Under Additional Properties, click **Web Services: Server Security Bindings**.
4. Under Request Receiver Binding, click **Edit > Encryption Information**.
5. Click **New** to create an encryption method.
6. Click **Delete** to delete an encryption method.

### **Encryption Information:**

Specifies the name of the encryption information.

## Encryption information configuration settings

Use this page to configure the encryption and decryption parameters. You can use these parameters to encrypt and decrypt various parts of the message including the body and user name token.

To view this administrative console page:

1. Click **Applications > Enterprise Applications** > *application\_name*.
2. Under Related Items, click **Web Module** > *URI\_file\_name* > **Web Services: Server Security Bindings**.

3. Under Request Receiver Binding, click **Edit > Encryption Information > New**.

**Encryption Information Name:**

Specifies the name for the encryption information.

**Key Locator Reference:**

Specifies the name used to reference the key locator.

To specify key locator references, click **Servers > Application Servers > *server\_name***. Under Additional Properties, click **Web Services: Default bindings for Web Services Security > Key Locators**.

**Encryption Key Name:**

Specifies the name of the encryption key, which is resolved to the actual key by the specified key locator.

**Key Encryption Algorithm:**

Specifies the algorithm Uniform Resource Identifier (URI) of the key encryption method.

The following algorithms are supported:

- [http://www.w3.org/2001/04/xmlenc#rsa-1\\_5](http://www.w3.org/2001/04/xmlenc#rsa-1_5)
- <http://www.w3.org/2001/04/xmlenc#kw-tripledes>

**5.1 +** The following additional algorithms are supported:

- <http://www.w3.org/2001/04/xmlenc#kw-aes128>
- <http://www.w3.org/2001/04/xmlenc#kw-aes256>
- <http://www.w3.org/2001/04/xmlenc#kw-aes192>

**Data Encryption Algorithm:**

Specifies the algorithm URI of the data encryption method.

The following algorithm is supported:

- <http://www.w3.org/2001/04/xmlenc#tripledes-cbc>

**5.1 +** The following additional algorithms are supported:

- <http://www.w3.org/2001/04/xmlenc#aes128-cbc>
- <http://www.w3.org/2001/04/xmlenc#aes256-cbc>
- <http://www.w3.org/2001/04/xmlenc#aes192-cbc>

## Encryption information configuration settings

Use this page to configure the encryption and decryption parameters.

The specifications listed on this page for the signature method, digest method, and canonicalization method are located in the World Wide Web Consortium (W3C) document entitled, *XML Encryption Syntax and Processing: W3C Recommendation 10 Dec 2002*.

To view this administrative console page, complete the following steps:

1. Click **Applications > Enterprise Applications > *application\_name***.
2. Under Related Items, click **Web Module > *URI\_file\_name* > Web Services: Server Security Bindings**.



3. Under Response Sender Binding, click **Edit > Encryption Information**.
4. If the encryption information is not available, select **None**.
5. If the encryption information is available, select **Dedicated Encryption Information**.

Then, specify the configuration in the following fields:

***Encryption Information Name:***

Specifies the name for the encryption information.

***Key Locator Reference:***

Specifies the name used to reference the key locator.

To specify key locator references, click **Servers > Application Servers > *server\_name***. Under Additional Properties, click **Web Services: Default bindings for Web Services Security > Key Locators**.

***Encryption Key Name:***

Specifies the name of the encryption key, which is resolved to the actual key by the specified key locator.

***Key Encryption Algorithm:***

Specifies the algorithm URI of the key encryption method.

The following algorithms are supported:

- [http://www.w3.org/2001/04/xmlenc#rsa-1\\_5](http://www.w3.org/2001/04/xmlenc#rsa-1_5)
- <http://www.w3.org/2001/04/xmlenc#kw-tripledes>

**5.1+** The following additional algorithms are supported:

- <http://www.w3.org/2001/04/xmlenc#kw-aes128>
- <http://www.w3.org/2001/04/xmlenc#kw-aes256>
- <http://www.w3.org/2001/04/xmlenc#kw-aes192>

By default the Java Cryptography Extension (JCE) is shipped with restricted or limited strength ciphers. To use 192-bit and 256-bit Advanced Encryption Standard (AES) encryption algorithms, you must apply unlimited jurisdiction policy files. Before downloading these policy files, back up the existing policy files (`local_policy.jar` and `US_export_policy.jar` in the `jre/lib/security/` directory) prior to overwriting them in case you want to restore the original files later. To download the policy files, complete either of the following sets of steps:

- For WebSphere Application Server platforms using IBM Developer Kit, Java Technology Edition Version 1.4.1, including the AIX, Linux, and Windows platforms, you can obtain unlimited jurisdiction policy files by completing the following steps:

- Go to the following Web site: <http://www.ibm.com/developerworks/java/jdk/security/index.html>
- Click **IBM SDK Policy files**.

The `unrestrict.zip` file is downloaded onto your machine.

- For WebSphere Application Server platforms using the Sun-based Java Development Kit (JDK) Version 1.4.1, including the Solaris environments and the HP-UX platform, you can obtain unlimited jurisdiction policy files by completing the following steps:

- Go to the following Web site: <http://java.sun.com/j2se/1.4.1/download.html>
- Locate the Java Cryptography Extension (JCE) Unlimited Strength Jurisdiction Policy Files 1.4.1 information and click **Download**. The `jce_policy-1_4_1.zip` file is downloaded onto your machine.

**Attention:**

After following either of these sets of steps, two Java Archive (JAR) files are placed in the JVM `jre/lib/security/` directory.

### **Data Encryption Algorithm:**

Specifies the algorithm Uniform Resource Identifiers (URI) of the data encryption method.

The following algorithm is supported:

- <http://www.w3.org/2001/04/xmlenc#tripleDES-cbc>

**5.1 +** The following additional algorithms are supported:

- <http://www.w3.org/2001/04/xmlenc#aes128-cbc>
- <http://www.w3.org/2001/04/xmlenc#aes256-cbc>
- <http://www.w3.org/2001/04/xmlenc#aes192-cbc>

By default the JCE is shipped with restricted or limited strength ciphers. To use 192-bit and 256-bit AES encryption algorithms, you must apply unlimited jurisdiction policy files.

### **Login bindings configuration settings**

Use this page to configure the encryption and decryption parameters.

The pluggable token uses the Java Authentication and Authorization Service (JAAS) `CallbackHandler` (`javax.security.auth.callback.CallbackHandler`) interface to generate the token that is inserted into the message. The following list describes the `Callback` support implementations:

#### **com.ibm.wsspi.wssecurity.auth.callback.BinaryTokenCallback**

This implementation is used for generating binary tokens inserted as `<wsse:BinarySecurityToken/@ValueType>` in the message.

#### **javax.security.auth.callback.NameCallback and javax.security.auth.callback.NameCallback**

This implementation is used for generating user name tokens inserted as `<wsse:UsernameToken>` in the message.

#### **com.ibm.wsspi.wssecurity.auth.callback.XMLTokenSenderCallback**

This implementation is used to generate Extensible Markup Language (XML) tokens and is inserted as the `<SAML: Assertion>` element in the message.

#### **com.ibm.wsspi.wssecurity.auth.callback.PropertyCallback**

This implementation is used to obtain properties specified in the binding file.

To view this administrative console page, complete the following steps:

1. Click **Applications > Enterprise Applications > application\_name**.
2. Under Related Items, click **Web Module > URI\_file\_name > Web Services: Client Security Bindings**.
3. Under Request Sender Bindings, click **Edit > Login Binding**.

If the encryption information is not available, select **None**.

If the encryption information is available, select **Dedicated Login Binding** and specify the configuration in the following fields:

### **Authentication Method:**

Specifies the unique name for the authentication method.

### **Callback Handler:**

Specifies the name of the callback handler. The callback handler must implement the `javax.security.auth.callback.CallbackHandler` interface.

**Basic Auth User ID:**

Specifies the user name for basic authentication. The basic authentication method provides the capability to define a user ID and a password in the binding file.

**Basic Auth Password:**

Specifies the password for basic authentication.

**Token Type URI:**

Specifies the Uniform Resource Identifiers (URI) for the token type. This information is inserted as `<wsse:BinarySecurityToken>/ValueType` for the `<SAML: Assertion>` XML token.

**Token Type Local Name:**

Specifies the local name for the token type. This information is inserted as `<wsse:BinarySecurityToken>/ValueType` for the `<SAML: Assertion>` XML token.

## Request sender

The security handler on the request sender side of the SOAP message enforces the security constraints, located in the `ibm-webservicesclient-ext.xml` file, and bindings, located in the `ibm-webservicesclient-bnd.xml` file. These constraints and bindings apply both to J2EE application clients or when Web services is acting as a client. The security handler acts on the security constraints before sending the SOAP message. For example, the security handler might digitally sign the message, encrypt the message, create a time stamp, or insert a security token.

The security handler on the request sender side of the Simple Object Access Protocol (SOAP) message enforces the security constraints, located in the `ibm-webservicesclient-ext.xml` file, and the bindings, located in the `ibm-webservicesclient-bnd.xml` file. These constraints and bindings apply both to J2EE application clients or when Web services is acting as a client. The security handler acts on the security constraints before sending the SOAP message. Request sender security constraints must match the security constraint requirements defined in the request receiver. For example, the security handler might digitally sign the message, encrypt the message, create a time stamp, or insert a security token. You can specify the following security requirements for the request sender and apply them to the SOAP message:

### Integrity (digital signature)

You can select multiple parts of a message to sign digitally. The following list contains the integrity options:

- Body
- Time stamp
- Security token

### Confidentiality (encryption)

You can select multiple parts of a message to encrypt. The following list contains the confidentiality options:

- Body content
- Username token

### Security token

You can insert only one token into the message. The following list contains the security token options:

- Basic authentication, which requires both a user name and a password

- Identity assertion, which requires a user name only
- X.509 binary security token
- Lightweight Third Party Authentication (LTPA) binary security token
- Custom token , which is pluggable and supports custom-defined tokens in the SOAP message

### Timestamp

You can have a time stamp to indicate the timeliness of the message.

- Timestamp

### ***Request sender binding collection:***

Use this page to specify the binding configuration to send request messages for Web services security.

To view this administrative console page, complete the following steps:

1. Click **Applications > Enterprise Applications > application\_name**.
2. Under Related Items, click **Web Modules > URI\_file\_name > Web Services: Client Security Bindings**.
3. Under Request Sender Binding, click **Edit**.

### *Signing Information:*

Specifies the configuration for the signing parameters. Signing information is used to sign and validate parts of the message including the body and time stamp.

You can also use these parameters for X.509 validation when the Authentication Method is `IDAssertion` and the *ID Type* is `X509Certificate` in the server-level configuration. In such cases, you must fill in the Certificate Path fields only.

### *Encryption Information:*

Specifies the configuration for the encrypting and decrypting parameters. Encryption information is used for encrypting and decrypting various parts of a message including the body and user name token.

### *Key Locators:*

Specifies a list of key locator objects that retrieve the keys for digital signature and encryption from a key store file or a repository. The key locator maps a name or a logical name to an alias or maps an authenticated identity to a key. This logical name is used to locate a key in a key locator implementation.

### *Login Mappings:*

Specifies a list of configurations for validating tokens within incoming messages.

Login mappings map the authentication method to the Java Authentication and Authorization Service (JAAS) configuration.

To configure JAAS, use the administrative console and click **Security > JAAS Configuration**.

## **Configuring the client for request encryption: Encrypting the message parts**

Prior to completing these steps, read either of the following topics to familiarize yourself with the Security Extensions tab and the Port Binding tab in the Web services client editor within the Assembly Toolkit:

- Configuring the client security bindings using the Assembly Toolkit
- Configuring the security bindings on a server acting as a client using the administrative console

These two tabs are used to configure the Web services security extensions and Web services security bindings, respectively.

Complete the following steps to specify which message parts to encrypt when configuring the client for request encryption:

1. Launch the Assembly Toolkit.
2. Open the J2EE perspective by clicking **Window > Open Perspective > Other > J2EE**.
3. Select the Web services-enabled Enterprise JavaBeans (EJB) or Web module.
4. In the Project Navigator window, locate the META-INF directory for an EJB module or the WEB-INF directory for a Web module.
5. Right-click the webservicessclient.xml file, and click **Open With > Web Services Client Editor**.
6. Click the **Security Extensions** tab, which is located at the bottom of the Web services client editor within the Assembly Toolkit.
7. Expand **Request Sender Configuration > Confidentiality**. Confidentiality refers to encryption while integrity refers to digital signing. Confidentiality reduces the risk of someone understanding the message flowing across the Internet. With confidentiality specifications, the message is encrypted before it is sent and decrypted when it is received at the correct target. For more information on encrypting , see XML encryption.
8. Select the parts of the message that you want to encrypt by clicking **Add**. You can select one of the following parts:

**Bodycontent**

User data portion of the message

**UsernameToken**

Basic authentication information, if selected

After you specify which message parts to encrypt, you must specify which method to use to encrypt the request message. See Configuring the client for request encryption: Choosing the encryption method for more information.

## Configuring the client for request encryption: Choosing the encryption method

Prior to completing these steps, read either of the following topics to familiarize yourself with the **Security Extensions** tab and the **Port Binding** tab in the Web Services Client Editor within the Assembly Toolkit:

- Configuring the client security bindings using the Assembly Toolkit
- Configuring the security bindings on a server acting as a client using the administrative console

These two tabs are used to configure the Web Services Security extensions and Web Services Security bindings, respectively.

Complete the following steps to specify which encryption method to use when configuring the client for request encryption:

1. Launch the Assembly Toolkit.
2. Open the J2EE perspective by clicking **Window > Open Perspective > Other > J2EE**.
3. Select the Web services-enabled Enterprise JavaBean (EJB) or Web module.
4. In the Project Navigator window, locate the META-INF directory for an EJB module or the WEB-INF directory for a Web module.
5. Right-click the webservicessclient.xml file, and click **Open With > Web Services Client Editor**.
6. Click the **Port Binding** tab, which is located at the bottom of the Web services client editor within the Assembly Toolkit.
7. Expand **Security Request Sender Binding Configuration > Encryption Information**.

8. Select an encryption option and click **Edit** to view the encryption information or click **Add** to add another option. The following table describes the purpose of this information. Some of these definitions are based on the XML-Encryption specification, which is located at the following Web address: <http://www.w3.org/TR/xmlenc-core>

**Encryption name**

Refers to the name of the encryption information entry.

**Data encryption method algorithm**

Encrypts and decrypts data in fixed size, multiple octet blocks.

**Key encryption method algorithm**

Represents public key encryption algorithms that are specified for encrypting and decrypting keys.

**Encryption key name**

Represents a Subject (Owner field of the certificate) from a public key certificate found by the encryption key locator, which is used by the key encryption method algorithm to encrypt the private key. The private key is used to encrypt the data.

The key chosen must be a public key of the target. Encryption must be done using the public key and decryption must be done by the target using the private key (the personal certificate of the target).

**Encryption key locator**

Represents a reference to a key locator implementation class that locates the correct key store where the alias and the certificate exist. For more information on configuring key locators, see *Configuring key locators using the Assembly Toolkit* and *Configuring key locators using the administrative console*.

For more information, see *Configuring key locators using the Assembly Toolkit* and *Configuring key locators using the Administrative Console*.

You must specify which parts of the request message to encrypt. See *Configuring the client for request encryption: Encrypting the message parts if you have not previously specified this information*.

**Request receiver**

The security handler on the request receiver side of the Simple Object Access Protocol (SOAP) message enforces the security specifications defined in the IBM extension deployment descriptor (`ibm-webservices-ext.xmi`) and bindings (`ibm-webservices-bnd.xmi`). The request receiver defines the security requirement of the SOAP message. The security constraint for request sender must match the security requirement of the request receiver for the server to accept the request. If the incoming SOAP message does not meet all the security requirements defined, then the request is rejected with the appropriate fault code returned to the sender. For security tokens, the token is validated using Java Authentication and Authorization Service (JAAS) login configuration and authenticated identity is set as the identity for the downstream invocation.

For example, if there is a security requirement to have the SOAP body digitally signed by Joe Smith and if the SOAP body of the incoming SOAP message is not signed by Joe Smith, then the request is rejected.

You can define the following security requirements for the request receiver:

**Required integrity (digital signature)**

You can select multiple parts of a message to sign digitally. The following list contains the integrity options:

- Body
- Time stamp
- Security token

### **Required confidentiality (encryption)**

You can select multiple parts of a message to encrypt. The following list contains the confidentiality options:

- Body content
- Token

You can have multiple security tokens. The following list contains the security token options:

- Basic authentication, which requires both a user name and a password
- Identity assertion, which requires a user name only
- X.509 binary security token
- Lightweight Third Party Authentication (LTPA) binary security token
- Custom token, which is pluggable and supports custom-defined tokens validated by the JAAS login configuration

### **Received time stamp**

You can have a time stamp for checking the timeliness of the message.

- Time stamp

### ***Request receiver binding collection:***

Use this page to specify the binding configuration to receive request messages for Web services security.

To view this administrative console page, complete the following steps:

1. Click **Applications > Enterprise Applications > application\_name**.
2. Under Related Items, click **Web Modules > URI\_file\_name > Web Services: Server Security Bindings**.
3. Under Request Receiver Binding, click **Edit**.

### ***Signing Information:***

Specifies the configuration for the signing parameters. Signing information is used to sign and validate parts of a message including the body, the time stamp, and the user name token.

You also can use these parameters for X.509 certificate validation when the authentication method is IDAssertion and the ID Type is X509Certificate in the server-level configuration. In such cases, you must fill in the Certificate Path fields only.

### ***Encryption Information:***

Specifies the configuration for the encrypting and decrypting parameters. This configuration is used to encrypt and decrypt parts of the message that include the body and the user name token.

### ***Trust Anchors:***

Specifies a list of key store objects that contain the trusted root certificates that are issued by a certificate authority (CA).

The certificate authority authenticates a user and issues a certificate. The CertPath API uses the certificate to validate the certificate chain of incoming, X.509-formatted security tokens or trusted, self-signed certificates.

### ***Collection Certificate Store:***

Specifies a list of the untrusted, intermediate certificate files.



The collection certificate store contains a chain of untrusted, intermediate certificates. The CertPath API attempts to validate these certificates, which are based on the trust anchor.

#### *Key Locators:*

Specifies a list of key locator objects that retrieve the keys for digital signature and encryption from a key store file or a repository. The key locator maps a name or a logical name to an alias or maps an authenticated identity to a key. This logical name is used to locate a key in a key locator implementation.

#### *Trusted ID Evaluators:*

Specifies a list of trusted ID evaluators that determine whether to trust the identity-asserting authority or message sender.

The trusted ID evaluators are used to authenticate additional identities from one server to another server. For example, a client sends the identity of user A to server 1 for authentication. Server 1 calls downstream to server 2, asserts the identity of user A, and includes the user ID and password of server 1. Server 2 attempts to establish trust with server 1 by authenticating its user ID and password and checking the trust based on the TrustedIDEvaluator implementation. If the authentication process and the trust check are successful, server 2 trusts that server 1 authenticated user A and a credential is created for user A on server 2 to invoke the request.

#### *Login Mappings:*

Specifies a list of configurations for validating tokens within incoming messages.

Login mappings map the authentication method to the Java Authentication and Authorization Service (JAAS) configuration.

To configure JAAS, use the administrative console and click **Security > JAAS Configuration**.

## **Configuring the server for request decryption: Decrypting the message parts**

Complete this task to specify which parts of the request message must be decrypted by the server. You must know which parts of the request message the client encrypts because the server must decrypt the same message parts.

Prior to completing these steps, read either of the following topics to become familiar with the Security Extensions tab and the Binding Configurations tab:

- Configuring the server security bindings using the Assembly Toolkit
- Configuring the server security bindings using the administrative console

These two tabs are used to configure the Web services security extensions and Web services security bindings, respectively.

1. Launch the Assembly Toolkit.
2. Open the J2EE perspective by clicking **Window > Open Perspective > Other > J2EE**.
3. Select the Web services-enabled Enterprise JavaBean (EJB) or Web module.
4. In the Project Navigator window, locate the META-INF directory for an EJB module or the WEB-INF directory for a Web module.
5. Right-click the webservices.xml file, and click **Open With > Web Services Editor**.
6. Click the **Security Extensions** tab, which is located at the bottom of the Web services editor within the Assembly Toolkit.
7. Expand the **Request Receiver Service Configuration Details > Required Confidentiality** section.

8. Select the parts of the message to decrypt. The message parts selected for the request decryption on the server must match the message parts selected for the message encryption on the client. Click **Add** and select either of the following message parts:

**bodycontent**

The user data section of the message.

**usertoken**

This token is the basic authentication information.

After you specify which parts of the request message to decrypt, you must specify the method to use to decrypt the message. See *Configuring the server for request decryption: Choosing the decryption method* for more information.

## Configuring the server for request decryption: choosing the decryption method

Prior to completing these steps, read either of the following topics to become familiar with the **Security Extensions** tab and the **Binding Configurations** tab:

- *Configuring the server security bindings using the Assembly Toolkit*
- *Configuring the server security bindings using the administrative console*

These two tabs are used to configure the Web Services Security extensions and Web Services Security bindings, respectively.

Complete this task to specify which decryption method is used by the server to decrypt the request message. You must know which decryption method the client uses because the server must use the same method.

1. Launch the Assembly Toolkit.
2. Open the J2EE perspective by clicking **Window > Open Perspective > Other > J2EE**.
3. Select the Web services-enabled Enterprise JavaBean (EJB) or Web module.
4. In the Project Navigator window, locate the META-INF directory for an EJB module or the WEB-INF directory for a Web module.
5. Right-click the webservices.xml file, select **Open With > Web Services Editor**.
6. Click the **Binding Configurations** tab, which is located at the bottom of the Web services editor within the Assembly Toolkit.
7. Expand the **Request Receiver Binding Configuration Details > Encryption Information** section.
8. Click **Edit** to view the encryption information. The following table describes the purpose for each of these selections. Some definitions are taken from the XML-Encryption specification, which is located at the following Web address: <http://www.w3.org/TR/xmlenc-core>

**Encryption name**

Represents the name of this encryption information entry; an alias for the entry.

**Data encryption method algorithm**

Encrypts and decrypts data in fixed size, multiple octet blocks. This algorithm must be the same as the algorithm selected in the client request sender configuration.

**Key encryption method algorithm**

Represents algorithms specified for encrypting and decrypting keys. This algorithm must be the same as the algorithm selected in the client request sender configuration.

**Encryption key name**

Represents a Subject from a personal certificate, which is typically a distinguished name (DN) that is found by the encryption key locator. The subject is used by the key encryption method algorithm to decrypt the secret key, and the secret key is used to decrypt the data.

The key chosen must be a private key in the key store configured by the key locator. The key requires the same Subject used by the client to encrypt the data. Encryption must be done using the public key and decryption by using the private key (personal certificate). To ensure that the client encrypts the data with the correct public or private key, you must extract the public key from the server key store and add it to the key store specified in the encryption configuration information for the client request sender.

For example, the personal certificate of a server is CN=Bob, O=IBM, C=US. Therefore the server contains the public and private key pair. The client sending the request should encrypt the data using the public key for CN=Bob, O=IBM, C=US. The server decrypts the data using the private key for CN=Bob, O=IBM, C=US.

### **Encryption key locator**

Represents a reference to a key locator implementation class that finds the correct keystore where the alias and the certificate exist. For more information on configuring key locators, go to the following sections: [Configuring key locators using the Assembly Toolkit](#) and [Configuring key locators using the Administrative Console](#).

It is important to note that for decryption, the encryption key name chosen must refer to a personal certificate that can be located by the key locator of the server referenced in the encryption information. Enter the Subject of the personal certificate here, which is typically a Distinguished Name (DN). The Subject uses the default key locator to find the key. If a custom key locator is written, the encryption key name can be anything used by the key locator to find the correct encryption key. The encryption key locator references the implementation class that finds the correct key store where this alias and certificate exist. Refer to [Configuring key locators using the Assembly Toolkit](#) and [Configuring key locators using the Administrative Console](#) for more information.

You must specify which parts of the request message to decrypt. See [Configuring the server for request decryption: Decrypting the message parts](#) if you have not previously specified this information.

### **Response sender**

The response sender defines the security requirements of the Simple Object Access Protocol (SOAP) response message. The security handler acts on the security constraints defined for the response in the IBM extension deployment descriptors, located in the `ibm-webservices-ext.xmi` file and the bindings, located in the `ibm-webservices-bnd.xmi` file. The security handler signs, encrypts, or generates the time stamp for the SOAP response message before the response is send to the caller.

#### **Integrity constraints (digital signature)**

You can select which parts of the message are digitally signed.

- Body
- Time stamp

#### **Confidentiality (encryption)**

You can encrypt the body content of the message.

#### **Time stamp**

You can have a time stamp for checking the timeliness of the message.

The security constraints that apply to the SOAP response message must match the security requirements defined in the response receiver. Otherwise, the response is rejected by the response receiver (caller).

#### ***Response sender binding collection:***

Use this page to specify the binding configuration for sender response messages for Web services security.

To view this administrative console page, complete the following steps:

1. Click **Applications > Enterprise Applications > *application\_name***.

2. Under Related Items, click **Web Modules** > *URI\_file\_name* > **Web Services: Server Security Bindings**.
3. Under Response Sender Binding, click **Edit**.

#### *Signing Information:*

Specifies the configuration for the signing parameters.

You also can use these parameters for X.509 certificate validation when the authentication method is IDAssertion and the ID Type is X509Certificate in the server-level configuration. In such cases, you must fill-in the **Certificate Path** fields only.

#### *Encryption Information:*

Specifies the configuration for the encryption and decryption parameters.

#### *Key Locators:*

Specifies a list of key locator objects that retrieve the keys for a digital signature and encryption from a key store file or a repository. The key locator maps a name or logical name to an alias or maps an authenticated identity to a key. This logical name is used to locate a key in a key locator implementation.

## **Configuring the server for response encryption: encrypting the message parts**

Prior to completing these steps, read either of the following topics to become familiar with the Security Extensions tab and the Binding Configurations tab in the Web services editor within the Assembly Toolkit:

- Configuring the server security bindings using the Assembly Toolkit
- Configuring the server security bindings using the administrative console

These two tabs are used to configure the Web services security extensions and the Web services security bindings, respectively.

Complete the following steps to specify which parts of the response message to encrypt when configuring the server for response encryption:

1. Launch the Assembly Toolkit.
2. Open the J2EE perspective by clicking **Window** > **Open Perspective** > **Other** > **J2EE**.
3. Select the Web services-enabled Enterprise JavaBean (EJB) or Web module.
4. In the Project Navigator window, locate the META-INF directory for an EJB module or the WEB-INF directory for a Web module.
5. Right-click the `webservices.xml` file, select **Open With** > **Web Services Editor**.
6. Click the **Security Extensions** tab, which is located at the bottom of the Web Services Editor within the Assembly Toolkit.
7. Expand **Response Sender Service Configuration Details** > **Confidentiality**. Confidentiality refers to encryption while integrity refers to digital signing. Confidentiality reduces the risk of someone understanding the message flowing across the Internet. With confidentiality specifications, the response is encrypted before it is sent and decrypted when it is received at the correct target. For more information on encrypting, see XML encryption.
8. Select the parts of the response that you want to encrypt by clicking **Add** and selecting **Bodytoken** or **UsernameToken**. The following information describes the message parts:

#### **Bodycontent**

User data portion of the message.

#### **UsernameToken**

Basic authentication information, if selected.

A user name token does not appear in the response so you do not need to select this option for the response. If you select this option, make sure that you also select it for the client response receiver. If you do not select this option, make sure that you do not select it for the client response receiver.

After you specify which message parts to encrypt, you must specify which method to use message encryption. See *Configuring the server for response encryption: Choosing the encryption method* for more information.

## Configuring the server for response encryption: Choosing the encryption method

Prior to completing these steps, read either of the following topics to become familiar with the Security Extensions tab and the Binding Configurations tab in the Web services editor within the Assembly Toolkit:

- *Configuring the server security bindings using the Assembly Toolkit*
- *Configuring the server security bindings using the administrative console*

These two tabs are used to configure the Web services security extensions and Web services security bindings, respectively.

Complete the following steps to specify which method the server uses to encrypt the response message:

1. Launch the Assembly Toolkit.
2. Click **Windows > Open Perspective > J2EE** to access the Assembly Toolkit perspective.
3. Select the Web services enabled Enterprise JavaBean (EJB) or Web module.
4. In the Project Navigator window, locate the META-INF directory for an EJB module or the WEB-INF directory for a Web module.
5. Right-click the `webservices.xml` file, and click **Open With > Web Services Editor**.
6. Click the **Binding Configurations** tab, which is located at the bottom of the Web Services Editor within the Assembly Toolkit.
7. Expand **Response Sender Binding Configuration Details > Encryption Information**.
8. Click **Edit** to view the encryption information. The following table describes the purpose of this information. Some of these definitions are based on the XML-Encryption specification, which is located at the following Web address: <http://www.w3.org/TR/xmlenc-core>

### Encryption name

Refers to the name of the encryption information entry.

### Data encryption method algorithm

Encrypts and decrypts data in fixed size, multiple octet blocks. The algorithm selected for the server response sender configuration must match the algorithm selected in the client response receiver configuration.

### Key encryption method algorithm

Represents public key encryption algorithms that are specified for encrypting and decrypting keys. The algorithm selected for the server response sender configuration must match the algorithm selected in the client response receiver configuration.

### Encryption key name

Represents a Subject from a public key certificate typically distinguished name (DN) that is found by the encryption key locator and used by the key encryption method algorithm to encrypt the private key. The private key is used to encrypt the data.

The key name chosen in the server response sender encryption information must be the public key of the key configured in the client response receiver encryption information. Encryption by the response sender must be done using the public key and decryption must be done by the response receiver using the associated private key (the personal certificate of the response receiver).

## Encryption key locator

The encryption key locator represents a reference to a key locator implementation class that finds the correct key store where the alias and the certificate exist. For more information on configuring key locators, see *Configuring key locators using the Assembly Toolkit* and *Configuring key locators using the administrative console*.

The encryption key name chosen must refer to a public key of the response receiver. For the encryption key name, use the Subject of the public key certificate, typically a Distinguished Name (DN). The name chosen is used by the default key locator to find the key. If you write a custom key locator, the encryption key name might be anything used by the key locator to find the correct encryption key (a public key). The encryption key locator references the implementation class that finds the correct key store where the alias and certificate exist. For more information, see *Configuring key locators using the Assembly Toolkit* and *Configuring key locators using the Administrative Console*.

You must specify which parts of the response message to encrypt. See *Configuring the server for response encryption: Encrypting the message parts* if you have not previously specified this information.

## Response receiver

The response receiver defines the security requirements of the response received from a request to a Web service. The security constraints for response sender must match the security requirements of the response receiver. If the constraints do not match, the response is not accepted by the caller or the sender. The security handler enforces the security constraints based on the security requirements defined in the IBM extension deployment descriptor, located in the `ibm-webservicesclient-ext.xml` file and in the bindings, located in the `ibm-webservicesclient-bnd.xml` file.

For example, the security requirement might have the response Simple Object Access Protocol (SOAP) body encrypted. If the SOAP body of the SOAP message is not encrypted, the response is rejected and the appropriate fault code is communicated back to the caller of the Web services.

You can specify the following security requirements for a response receiver:

### Required integrity (digital signature)

You can select which parts of a message are digitally signed. The following list contains the integrity options:

- Body
- Time stamp

### Required confidentiality (encryption)

You can encrypt the body content of the message.

### Received time stamp

You can have a time stamp for checking the timeliness of the message.

### ***Response receiver binding collection:***

Use this page to specify the binding configuration for receiver response messages for Web services security.

To view this administrative console page, complete the following steps:

1. Click **Applications > Enterprise Applications > application\_name**.
2. Under Related Items, click **Web Modules > URI\_file\_name > Web Services: Server Security Bindings**.
3. Under Response Sender Binding, click **Edit**.

*Signing Information:*



Specifies the configuration for the signing parameters. Signing information is used to sign and to validate parts of the message including the body and the time stamp.

You can also use these parameters for X.509 validation when the authentication method is `IDAssertion` and the *ID type* is `X509Certificate` in the server-level configuration. In such cases, you must fill in the certificate path fields only.

#### *Encryption Information:*

Specifies the configuration for the encryption and decryption parameters. Encryption information is used for encrypting and decrypting various parts of a message, including the body and the user name token.

#### *Trust Anchors:*

Specifies a list of key store objects that contain the trusted root certificates that are self-signed or issued by a certificate authority.

The certificate authority authenticates a user and issues a certificate. After the certificate is issued, the key store objects, which contain these certificates, use the certificate for certificate path or certificate chain validation of incoming X.509-formatted security tokens.

#### *Collection Certificate Store:*

Specifies a list of the untrusted, intermediate certificate files.

The collection certificate store contains a chain of untrusted, intermediate certificates. The CertPath API attempts to validate these certificates, which are based on the trust anchor.

#### *Key Locators:*

Specifies a list of key locator objects that retrieve the keys for a digital signature and encryption from a key store file or a repository. The key locator maps a name or logical name to an alias or maps an authenticated identity to a key. This logical name is used to locate a key in a key locator implementation.

## **Configuring the client for response decryption: Decrypting the message parts**

Prior to completing these steps, read either of the following topics to become familiar with the Security Extensions tab and the Port Binding tab in the Web Services Client Editor within the Assembly Toolkit:

- Configuring the client security bindings using the Assembly Toolkit
- Configuring the security bindings on a server acting as a client using the administrative console

These two tabs are used to configure the Web services security extensions and the Web services security bindings, respectively.

Complete the following steps to specify which response message parts to decrypt when configuring the client for response decryption. The server response encryption and client response decryption configurations must match.

1. Launch the Assembly Toolkit.
2. Open the J2EE perspective by clicking **Window > Open Perspective > Other > J2EE**.
3. Select the Web services enabled EJB or Web module.
4. In the Project Navigator window, locate the META-INF directory for an EJB module or the WEB-INF directory for a Web module.
5. Right-click the `webservicesclient.xml` file, click here **Open With > Web Services Client Editor**.
6. Click the **Security Extensions** tab, which is located at the bottom of the Web services client editor within the Assembly Toolkit.



7. Expand the **Response Receiver Configuration > Required Confidentiality** section.
8. Select the parts of the message that you must decrypt by clicking **Add** and selecting either **Bodycontent** or **Usenametoken**. The following information describes these message parts:

**Bodycontent**

The user data portion of the message.

**Usenametoken**

The basic authentication information, if selected.

The information selected in this step is encrypted by the server in the response sender.

**Important:** A username token is typically not sent in the response. Thus, you usually do not need to select username token.

After you specify which message parts to decrypt, you must specify which method to use when decrypting the response message. See *Configuring the client for response decryption: Choosing a decryption method* for more information.

## Configuring the client for response decryption: Choosing a decryption method

Prior to completing these steps, read either of the following topics to become familiar with the **Security Extensions** tab and the **Port Binding** tab in the Web Services Client Editor within the Assembly Toolkit:

- *Configuring the client security bindings using the Assembly Toolkit*
- *Configuring the security bindings on a server acting as a client using the administrative console*

These two tabs are used to configure the Web services security extensions and Web services security bindings, respectively.

Complete the following steps to specify which decryption method to use when the client decrypts the response message. The server response encryption and client response decryption configurations must match.

1. Launch the Assembly Toolkit.
2. Open the J2EE perspective by clicking **Window > Open Perspective > Other > J2EE**.
3. Select the Web services enabled Enterprise JavaBean (EJB) or Web module.
4. In the Project Navigator window, locate the META-INF directory for an EJB module or the WEB-INF directory for a Web module.
5. Right-click the `webservicesclient.xml` file, and click **Open With > Web Services Client Editor**.
6. Click the **Port Binding** tab, which is located at the bottom of the Web services client editor within the Assembly Toolkit.
7. Expand the **Security Response Receiver Binding Configuration > Encryption Information** section. For more information on encrypting and decrypting Simple Object Access Protocol (SOAP) messages, see XML encryption.
8. Click **Edit** to view the encryption information. The following table describes the purpose for this information. Some of these definitions are based on the XML-Encryption specification, which is located at the following Web address: <http://www.w3.org/TR/xmlenc-core>

**Encryption name**

Refers to the alias used for the encryption information entry.

**Data encryption method algorithm**

Encrypts and decrypts data in fixed size, multiple octet blocks.

**Key encryption method algorithm**

Represents public key encryption algorithms specified for encrypting and decrypting keys.

### Encryption key name

Represents a Subject from a personal certificate, which is typically a distinguished name (DN) that is found by the encryption key locator. The Subject is used by the key encryption method algorithm to decrypt the secret key. The secret key is used to decrypt the data.

**Important:** The key chosen must be a private key of the client. Encryption must be done using the public key and decryption must be done by the private key (personal certificate). For example, the personal certificate of the client is: CN=Alice, O=IBM, C=US. Therefore, the client contains the public and private key pair. The target server that sends the response encrypts the secret key using the public key for CN=Alice, O=IBM, C=US. The client decrypts the secret key using the private key for CN=Alice, O=IBM, C=US

### Encryption key locator

The encryption key locator represents a reference to a key locator implementation class that finds the correct key store where the alias and the certificate exist. For more information on configuring key locators, see [Configuring key locators using the Assembly Toolkit](#) and [Configuring key locators using the administrative console](#).

For decryption, the encryption key name chosen must refer to a personal certificate that can be located by the client key locator. The Subject (owner field of the certificate) of the personal certificate should be entered in the Encryption key name, this is typically a Distinguished Name (DN). The default key locator uses the Encryption key name to find the key within the keystore. If you write a custom key locator, the encryption key name can be anything used by the key locator to find the correct encryption key. The encryption key locator references the implementation class that locates the correct key store where this alias and certificate exists. For more information, see [Configuring key locators using the Assembly Toolkit](#) and [Configuring key locators using the administrative console](#).

You must specify which parts of the request message to decrypt. See the topic [Configuring the client for response decryption: Decrypting the message parts if you have not previously specified this information](#).

## Securing Web services using basicauth authentication

WebSphere Application Server provides several different methods to secure your Web services. Extensible Markup Language (XML) digital signature is one of these methods. You might also secure your Web services using any of the following methods:

- XML digital signature
- XML encryption
- Basicauth authentication
- Identity assertion authentication
- Signature authentication
- Pluggable token

With the basicauth authentication method, the request sender generates a basicauth security token using a callback handler. The request receiver retrieves the basicauth security token from the SOAP message and validates it using a Java Authentication and Authorization Service (JAAS) login module. Trust is established using user name and password validation. To use basicauth authentication to secure Web services, complete the following tasks:

1. Secure the client for basicauth authentication.
  - a. Configure the client for basicauth authentication: Specifying the method
  - b. Configure the client for basicauth authentication: Collecting the authentication information
2. Secure the server for basicauth authentication.

- a. Configure the server to handle basicauth authentication
- b. Configure the server to validate basicauth authentication information

After completing these steps, you have secured your Web services using basicauth authentication.

## Configuring the client for basic authentication: Specifying the method

BasicAuth refers to the user ID and password of a valid user in the registry of the target server. Collection of BasicAuth information can occur in many ways including through a GUI prompt, a standard in (Stdin) prompt, or specified in the bindings, which prevents user interaction. For more information on BasicAuth authentication, see: BasicAuth authentication method.

**Attention:** **5.1+** WebSphere Application Server supports nonce (randomly generated token) with BasicAuth authentication. For more information, see Nonce.

Complete the following steps to specify BasicAuth as the authentication method:

1. Launch the Assembly Toolkit.
2. Open the J2EE perspective by clicking **Window > Open Perspective > Other > J2EE**.
3. Select the Web services enabled Enterprise JavaBean (EJB) module or Web module.
4. In the Project Navigator window, locate the META-INF directory for an EJB module or the WEB-INF directory for a Web module.
5. Right-click the webservicescClient.xml file, and click **Open With > Web Services Client Editor**.
6. Click the **Security Extensions** tab, which is located at the bottom of the Web services client editor within the Assembly Toolkit.
7. Expand the **Request Sender Configuration > Login Configuration** section. The only valid login configuration choices for a pure client are BasicAuth and Signature.
8. Select **BasicAuth** to authenticate the client using a user ID and a password. This user ID and password must be specified in the target user registry. The other choice, Signature, attempts to authenticate the client using the certificate used to digitally sign the message.

For more information on getting started with the Web services client editor within the Assembly Toolkit, see either of the following topics:

- Configuring the client security bindings using the Assembly Toolkit
- Configuring the security bindings on a server acting as a client using the administrative console

After you specify the BasicAuth authentication method, you must specify how to collect the authentication information. See Configuring the client for basic authentication: Collecting the authentication information.

### **BasicAuth authentication method:**

When you use the BasicAuth authentication method, the security token that is generated is a <wsse:UsernameToken> element with <wsse:Username> and <wsse>Password> elements.

WebSphere Application Server supports text passwords but not password digest because passwords are not stored and cannot be retrieved from the server. On the request sender side, a callback handler is invoked to generate the security token. On the request receiver side, a Java Authentication and Authorization Service (JAAS) login module is used to validate the security token. These two operations, token generation and token validation, are described in the following sections.

### **BasicAuth token generation**

The request sender generates a BasicAuth security token using a callback handler. The security token returned by the callback handler is inserted in the Simple Object Access Protocol (SOAP) message. The callback handler that is used is specified in the <LoginBinding> element of the

bindings file, `ibm-webservicesclient-bnd.xml`. The following callback handler implementations are provided with WebSphere Application Server and can be used with the BasicAuth authentication method:

- `com.ibm.wsspi.wssecurity.auth.callback.GUIPromptCallbackHandler`
- `com.ibm.wsspi.wssecurity.auth.callback.StdinPromptCallbackHandler`
- `com.ibm.wsspi.wssecurity.auth.callback.NonPromptCallbackHandler`

You can add your own callback handlers that implement the `javax.security.auth.callback.CallbackHandler` method.

### BasicAuth token validation

The request receiver retrieves the BasicAuth security token from the SOAP message and validates it using a JAAS login module. The `<wsse:Username>` and `<wsse:Password>` elements in the security token are used to perform the validation. If the validation is successful, the login module returns a JAAS Subject. This Subject is set as the identity of the running thread. If the validation fails, the request is rejected with a SOAP fault exception.

The JAAS login configuration is specified in the `<LoginMapping>` element of the bindings file. Default bindings are specified in the `ws-security.xml` file. However, you can override these bindings using the application-specific `ibm-webservices-bnd.xml` file. The configuration information consists of a `CallbackHandlerFactory` and a `ConfigName` value. The `CallbackHandlerFactory` option specifies the name of a class that is used for creating the JAAS `CallbackHandler` object. WebSphere Application Server provides the `com.ibm.wsspi.wssecurity.auth.callback.WSCallbackHandlerFactoryImpl` `CallbackHandlerFactory` implementation. The `ConfigName` value specifies a JAAS configuration name entry. WebSphere Application Server searches the `security.xml` file for a matching configuration name entry. If a match is not found, it searches the `wsjaas.conf` file for a match. WebSphere Application Server provides the `WSLogin` default configuration entry, which is suitable for the BasicAuth authentication method.

### Configuring the client for basic authentication: collecting the authentication information

BasicAuth refers to the user ID and the password of a valid user in the registry of the target server. Collection of BasicAuth information can occur in many ways including through a user interface prompt, a standard in (Stdin) prompt, or specified in the bindings, which prevents user interaction. For more information on BasicAuth authentication, see BasicAuth authentication method.

Complete this task to specify the authentication information needed for BasicAuth authentication:

1. Launch the Assembly Toolkit.
2. Open the J2EE perspective by clicking **Window > Open Perspective > Other > J2EE**.
3. Select the Web services enabled Enterprise JavaBean (EJB) or Web module.
4. In the Project Navigator window, locate the META-INF directory for an EJB module or the WEB-INF directory for a Web module.
5. Right-click the `webservicesclient.xml` file, and click **Open With > Web Services Client Editor**.
6. Click the **Port Binding** tab, which is located at the bottom of the Web services client editor within the Assembly Toolkit.
7. Expand the **Security Request Sender Binding Configuration > Login Binding** section.
8. Click **Edit** or **Enable** to view the login binding information. The login binding information displays and enter the following information:

#### Authentication method

Specifies the type of authentication. Select **BasicAuth** to use basic authentication.

### **Token value type URI and Token value type local name**

When you select BasicAuth, you cannot edit the token value type URI and the local name values. Specifies values for custom authentication types. For BasicAuth authentication, leave these values blank.

### **Callback handler**

Specifies the Java Authentication and Authorization Server (JAAS) callback handler implementation for collecting the BasicAuth information. You can use the following default implementations for the callback handler:

#### **com.ibm.wsspi.wssecurity.auth.callback.StdinPromptCallbackHandler**

This implementation is used for non-user interface console prompts.

#### **com.ibm.wsspi.wssecurity.auth.callback.GUIPromptCallbackHandler**

This implementation is used for user interface panel prompts.

#### **com.ibm.wsspi.wssecurity.auth.callback.NonPromptCallbackHandler**

This implementation is used when you plan to always enter the user ID and password in the BasicAuth user ID and password section that follows.

### **Basic Authentication user ID and Basic Authentication password**

Specifies values for the BasicAuth user ID and password, regardless of the default callback handler indicated previously, these user ID and password values are used to authenticate to the server for the Web services security authentication. If you leave these values blank, use either the GUIPromptCallbackHandler or the StdinPromptCallbackHandler implementation, but only on a pure client. Always fill-in these values for any Web service that acts as a client to another Web service that you want to specify for BasicAuth for authentication downstream. If you want the client identity of the originator to flow downstream, configure the Web service client to use either ID assertion or Lightweight Third Party Authentication (LTPA).

### **Property**

Specifies properties with name and value pairs for custom callback handlers to use. For BasicAuth authentication, you do not need to enter any information. To enter a new property, click **Add** and enter the new property and value.

**Other basic authentication entries:** There is a basic authentication entry in the Port Qualified Name Binding Details section. This entry is used for HTTP transport authentication, which might be required if the router servlet is protected.

Information specified in the Web services security basic authentication section overrides the basic authentication information specified in the Port Qualified Name Binding Details section for authorizing the Web service.

For a server that acts as a client, do not specify a user interface or non-user interface prompt callback handler. To configure BasicAuth authentication from one Web service to a downstream Web service, select the `com.ibm.wsspi.wssecurity.auth.callback.NonPromptCallbackHandler` implementation and explicitly specify the BasicAuth user ID and password. If you want the client identity of the originator to flow downstream, configure the Web service client to use ID assertion.

To use the BasicAuth authentication method, you must specify the method in the **Login Configuration** section of the Assembly Toolkit. See *Configuring the client for basicauth authentication: Specifying the method* if you have not previously specified this information.

### ***Identity assertion authentication method:***

When using the identity assertion (IDAssertion) authentication method, the security token generated is a `<wsse:UsernameToken>` element that contains a `<wsse:Username>` element.

On the request sender side, a callback handler is invoked to generate the security token. On the request receiver side, the security token is validated. These two operations, token generation and token validation operations, are described in the following sections.

#### **Identity assertion token validation:**

The request receiver retrieves the IDAssertion security token from the SOAP message and validates it using a Java Authentication and Authorization Service (JAAS) login module. With identity assertion, special processing is required to establish trust before asserting the identity as the established identity of the running thread. This special processing is defined by the <IDAssertion> element in the deployment descriptor file, `ibm-webservices-ext.xmi`. If all the validation checks are successful, the asserted identity is set as the identity of the running thread of. If the validation fails, the request is rejected with a SOAP fault exception.

The JAAS login configuration is specified in the <LoginMapping> element of the bindings file. Default bindings are specified in the `ws-security.xml` file. However, you can override these bindings using the application specific `ibm-webservices-bnd.xmi` file. The configuration information consists of `CallbackHandlerFactory` and a `ConfigName`. `CallbackHandlerFactory` specifies the name of a class that is used for creating the JAAS `CallbackHandler` object. WebSphere Application Server provides the `com.ibm.wsspi.wssecurity.auth.callback.WSCallbackHandlerFactoryImpl` `CallbackHandlerFactory` implementation. `ConfigName` specifies a JAAS configuration name entry.

WebSphere Application Server searches the `security.xml` file for a matching configuration name entry. If a match is not found it searches the `wsjaas.conf` file. WebSphere Application Server provides the `system.wssecurity.IDAssertion` default configuration entry, which is suitable for the identity assertion authentication method.

The <IDAssertion> element in the `ibm-webservices-ext.xmi` deployment descriptor file specifies the special processing required when using the identity assertion authentication method. The <IDAssertion> element is composed of two sub-elements: <IDType> and <TrustMode>.

The <IDType> element specifies the method for asserting the identity. The supported values for asserting the identity are:

- Username
- Distinguished name (DN)
- X.509 certificate

When <IDType> is *username*, a username token (for example, Bob) is provided. This user name is mapped to a user in the user registry and is the asserted identity after successful trust validation. When the <IDType> value is *DN*, a user name token containing a distinguished name is provided (for example, `cn=Bob Smith, o=ibm, c=us`). This DN is mapped to a user in the user registry and this user is the asserted identity after successful trust validation. When the <IDType> is *X509Certificate*, a binary security token containing an X509 certificate is provided and the SubjectDN value from the certificate (for example, `cn=Bob Smith, o=ibm, c=us`) is extracted. This SubjectDN value is mapped to a user in the user registry and this user is the asserted identity after successful trust validation.

The <TrustMode> element specifies how the trust authority, or asserting authority, provides trust information. The supported values are:

- Signature
- BasicAuth
- No value specified

When the <TrustMode> value is *Signature* the signature is validated. Then, the signer (for example, `cn=IBM Authority, o=ibm, c=us`) is mapped to an identity in the user registry (for example, `IBMAuthority`).



To ensure that the asserting authority is trusted, the mapped identity (for example, `IBMAuthority`) is validated against a list of trusted identities. When the `<TrustMode>` element is `BasicAuth`, there is a user name token with a user name and password, which is the user name and password of the asserting authority.

The user name and password are validated. If they are successfully validated, that user name (for example, `IBMAuthority`) is validated against a list of trusted identities. If a value is not specified for `<TrustMode>`, trust is presumed and additional trust validation is not performed. This type of identity assertion is called *presumed trust mode*. Use the presumed trust mode only in an environment where the trust is established using some other mechanism.

If all the validations described previously succeed, the asserted identity (for example, `Bob`) is set as the identity of the running thread. If any of the validations fail, the request is rejected with a SOAP fault exception.

## Configuring the server to handle `BasicAuth` authentication information

`BasicAuth` refers to the user ID and the password of a valid user in the registry of the target server. After a request is received that contains basic authentication information, the server needs to log in to form a credential. The credential is used for authorization. If the user ID and the password supplied are not valid, an exception is thrown and the request ends without invoking the resource. For more information on `BasicAuth` authentication, see `BasicAuth`.

Complete the following steps to configure the server to handle `BasicAuth` authentication information:

1. Launch the Assembly Toolkit.
2. Open the J2EE perspective by clicking **Window > Open Perspective > Other > J2EE**.
3. Select the Web services-enabled Enterprise JavaBean (EJB) or Web module.
4. In the Project Navigator window, locate the `META-INF` directory for an EJB module or the `WEB-INF` directory for a Web module.
5. Right-click the `webservices.xml` file, and click **Open With > Web Services Editor**.
6. Click the **Security Extensions** tab, which is located at the bottom of the Web services editor within the Assembly Toolkit.
7. Expand the **Request Receiver Service Configuration Details > Login Configuration** section. You can select the following options:
  - `BasicAuth`
  - Signature
  - ID assertion
  - Lightweight Third Party Authentication (LTPA)
8. Select **BasicAuth** to authenticate the client with a user ID and a password. The client must specify a valid user ID and password in the server user registry.

You can select multiple login configurations, which means that different types of security information might be received at the server. The order in which the login configurations are added decides the order in which they are processed when a request is received. Problems can occur if you have multiple login configurations added that have security tokens in common. For example, ID assertion contains a `BasicAuth` token. For ID assertion to work properly, list ID assertion ahead of `BasicAuth` in the processing list or the `BasicAuth` processing overrides the `IDAssertion` processing.

After you specify how the server handles `BasicAuth` authentication information, you must specify how the server validates the authentication information. See `Configuring the server to validate basicauth` authentication information for more information.



## Configuring the server to validate BasicAuth authentication information

BasicAuth refers to the user ID and the password of a valid user in the registry of the target server. Once a request is received that contains basic authentication information, the server needs to log in to form a credential. The credential is used for authorization. If the user ID and the password supplied is invalid, an exception is thrown and the request ends without invoking the resource. For more information on BasicAuth authentication, see BasicAuth.

Complete the following steps to specify how the server validates the BasicAuth authentication information:

1. Launch the Assembly Toolkit.
2. Open the J2EE perspective by clicking **Window > Open Perspective > Other > J2EE**.
3. Select the Web services-enabled Enterprise JavaBean (EJB) or Web module.
4. In the Project Navigator window, locate the META-INF directory for an EJB module or the WEB-INF directory for a Web module.
5. Right-click the `webservices.xml` file, select **Open With > Web Services Editor**.
6. Click the **Binding Configurations** tab, which is located at the bottom of the Web services editor within the Assembly Toolkit.
7. Expand the **Request Receiver Binding Configuration Details > Login Mapping** section.
8. Click **Edit** to view the login mapping information or click **Add** to add new login mapping information. The login mapping dialog is displayed. Select or enter the following information:

### Authentication method

Specifies the type of authentication that occurs. Select **BasicAuth** to use basic authentication.

### Configuration name

Specifies the Java Authentication and Authorization Service (JAAS) login configuration name. For the BasicAuth authentication method, enter `WSLogin` for the JAAS login Configuration name.

### Use token valid type

Determines if you want to specify a custom token type. For the default authentication method selections, you do not need to specify this option.

### Token value type URI and Token value type URI local name

When you select BasicAuth, you cannot edit the token value type URI and local name values. Specifies custom authentication types. For BasicAuth authentication leave these fields blank.

### Callback handler factory class name

Creates a JAAS CallbackHandler implementation that understands the following callbacks:

- `javax.security.auth.callback.NameCallback`
- `javax.security.auth.callback.PasswordCallback`
- `com.ibm.wsspi.wssecurity.auth.callback.BinaryTokenCallback`
- `com.ibm.wsspi.wssecurity.auth.callback.XMLTokenReceiverCallback`
- `com.ibm.wsspi.wssecurity.auth.callback.PropertyCallback`

### Callback handler factory property name and Callback handler factory property value

Specifies callback handler properties for custom callback handler factory implementations. You do not need to specify any properties for the default callback handler factory implementation. For BasicAuth, you do not need to enter any property values.

### Login mapping property name and Login mapping property value

Specifies properties for a custom login mapping. For the default implementations including BasicAuth, leave these fields blank.

You must specify how the server handles the BasicAuth authentication method. See Configuring the server to handle BasicAuth authentication information if you have not previously specified this information.

## Identity assertion

Identity assertion is a method for expressing the identity of the sender (for example, user name) in a SOAP message. When identity assertion is used as an authentication method, the authentication decision is performed based only on the name of the identity, and on other information such as passwords and certificates.

### ID type

The Web Services Security implementation in WebSphere Application Server can handle these identity types:

#### User name

Denotes the user name, such as the one in the local operating system (for example, "alice"). This name is embedded in the <Username> element within the <UsernameToken> element.

**DN** Denotes the distinguished name (DN) for the user, such as "CN=alice, O=IBM, C=US". This name is embedded in the <Username> element within the <UsernameToken> element.

#### X.509 certificate

Represents the identity of the user as an X.509 certificate instead of a string name. This certificate is embedded in the <BinarySecurityToken> element.

### Managing trust

The intermediary host in the SOAP message itinerary can assert claimed identity of the initial sender. Two methods (called *trust mode*) are supported for this assertion:

#### Basic authentication

The intermediary adds its user name and password pair to the message.

#### Signature

The intermediary digitally signs the <UsernameToken> element of the initial sender.

**Note:** This trust mode does not support the X.509 certificate ID type.

### Typical scenario

ID assertion is typically used in the multihop environment where the SOAP message passes through one or more intermediary hosts. The intermediary host authenticates the initial sender. The following scenario describes the process:

1. The initial sender sends a SOAP message to the intermediary host with some embedded authentication information. This authentication information might be a user name and a password pair with an Lightweight Third Party Authentication (LTPA) token.
2. The intermediary host authenticates the initial sender according to the embedded authentication information.
3. The intermediary host removes the authentication information from the SOAP message and replaces it with the <UsernameToken> element, which contains a user name.
4. The intermediary host asserts the trust according to the trust mode.
5. The intermediary host sends the updated SOAP message to the ultimate receiver.
6. The ultimate receiver checks the trust against the intermediary host information according to the configured trust mode. Also, the trusted ID evaluator is invoked.
7. If trust is established by the final receiver, the receiver invokes the Web service under the authorization of the user name (that is, the initial sender) in the SOAP message.

## Securing Web services using identity assertion authentication

WebSphere Application Server provides several different methods to secure your Web services. Extensible Markup Language (XML) digital signature is one of these methods. You might also secure your Web services using any of the following methods:

- XML digital signature
- XML encryption
- Basicauth authentication
- Identity assertion authentication
- Signature authentication
- Pluggable token

With the identity assertion authentication method, the security token generates a `<wsee:Username Token>` element that contains a `<wsse:Username>` element. On the request sender side, a callback handler is invoked to generate the security token. On the request receiver side, the security token is validated. Unlike basicauth authentication, trust is established through the use of a security token rather than through user name and password validation. To use identity assertion authentication to secure Web services, complete the following tasks:

1. Secure the client for identity assertion authentication.
  - a. Configure the client for identity assertion authentication: Specifying the method.
  - b. Configure the client for identity assertion authentication: Collecting the authentication information.
2. Secure the server for identity assertion authentication.
  - a. Configure the server to handle identity assertion authentication.
  - b. Configure the server to validate identity assertion authentication information.

After completing these steps, you have secured your Web services using identity assertion authentication.

### Configuring the client for identity assertion: specifying the method

This task is used to configure identity assertion authentication. The purpose of identity assertion is to assert the authenticated identity of the originating client from a Web service to a downstream Web service. Do not attempt to configure identity assertion from a pure client. Identity assertion works only when you configure on the client-side of a Web service acting as a client to a downstream Web service.

In order for the downstream Web service to accept the identity of the originating client (just the user name), you must supply a special trusted BasicAuth credential that the downstream Web service trusts and can authenticate successfully. You must specify the user ID of the special BasicAuth credential in a trusted ID evaluator on the downstream Web service configuration. For more information on trusted ID evaluators, see Trusted ID evaluators.

Complete the following steps to specify identity assertion as the authentication method:

1. Launch the Assembly Toolkit.
2. Open the J2EE perspective by clicking **Window > Open Perspective > Other > J2EE**.
3. Select the Web services-enabled Enterprise JavaBean (EJB) or Web module.
4. In the Project Navigator window, locate the META-INF directory for an EJB module or the WEB-INF directory for a Web module.
5. Right-click the `webservicesclient.xml` file, and click **Open With > Web Services Client Editor**.
6. Click the **Security Extensions** tab, which is located at the bottom of the Web services client editor within the Assembly Toolkit.
7. Expand the **Request Sender Configuration > Login Configuration** section.

8. Select **IDAssertion** as the authentication method. For more conceptual information on identity assertion authentication, see ID assertion.
9. Expand the **IDAssertion** section.
10. For the ID Type, select **Username**. This value works with all registry types and originating authentication methods.
11. For the Trust Mode, select either **BasicAuth** or **Signature**.
  - By selecting **BasicAuth**, you must include basic authentication information (user ID and password), which the downstream Web service has specified in the trusted ID evaluator as a trusted user ID. See Configuring the client for signature authentication: Collecting the authentication information to specify the user ID and password information.
  - By selecting **Signature** the certificate configured in the signature information section used to sign the data also is that is used as the trusted subject. The Signature is used to create a credential and user ID, which the certificate mapped to the downstream registry, is used in the trusted ID evaluator as a trusted user ID.

See Configuring the client security bindings using the Assembly Toolkit for more information on the Web services client editor within the Assembly Toolkit.

After you specify identity assertion as the authentication method used by the client, you must specify how to collect the authentication information. See Configuring the client for identity assertion authentication: Collecting the authentication information for more information.

## Configuring the client for identity assertion: Collecting the authentication method

This task is used to configure identity assertion authentication. The purpose of identity assertion is to assert the authenticated identity of the originating client from a Web service to a downstream Web service. Do not attempt to configure identity assertion from a pure client. Identity assertion works only when you configure on the client-side of a Web service acting as a client to a downstream Web service.

In order for the downstream Web service to accept the identity of the originating client (just the user name), you must supply a special trusted BasicAuth credential that the downstream Web service trusts and can authenticate successfully. You must specify the user ID of the special BasicAuth credential in a trusted ID evaluator on the downstream Web service configuration. For more information on trusted ID evaluators, see Trusted ID evaluators.

Complete the following steps to specify how the client collects the authentication information:

1. Launch the Assembly Toolkit.
2. Open the J2EE perspective by clicking **Window > Open Perspective > Other > J2EE**.
3. Select the Web services-enabled Enterprise JavaBean (EJB) or Web module.
4. In the Project Navigator window, locate the META-INF directory for an EJB module or the WEB-INF directory for a Web module.
5. Right-click the webservicessclient.xml file, and click **Open With > Web Services Client Editor**.
6. Click the **Port Binding** tab, which is located at the bottom of the Web Services Client Editor within the Assembly Toolkit.
7. Expand the **Security Request Sender Binding Configuration > Login Binding** section.
8. Click **Edit** to view the login binding information and select **IDAssertion**. The login binding dialog is displayed. Select or enter the following information:

### Authentication method

The authentication method specifies the type of authentication that occurs. Select **IDAssertion** to use identity assertion.

**Token value type URI and Token value type Local name**

When you select IDAssertion, you cannot edit the token value type Universal Resource Identifier (URI) and the local name. Specifies custom authentication types. For IDAssertion authentication, leave these values blank.

**Callback handler**

Specifies the Java Authentication and Authorization Service (JAAS) callback handler implementation for collecting the BasicAuth information. Specify the `com.ibm.wsspi.wssecurity.auth.callback.NonPromptCallbackHandler` implementation for IDAssertion.

**Basic authentication User ID and Basic authentication Password**

In this field the trust mode entered in the extensions is BasicAuth. Specifies the trusted user ID and password in these fields. The user ID specified must be an ID that is trusted by the downstream Web service. The Web service trusts the user ID if it is entered as a trusted ID in a trusted ID evaluator in the downstream Web service bindings. If the trust mode entered in the extensions is Signature, you do not need to specify any information in this field.

**Property Name and Property Value**

Specifies properties with name and value pairs, for use by custom callback handlers. For IDAssertion, you do not need to specify any information in this field.

To use the identity assertion authentication method, you must specify the method in the **Security Extensions** section of the Assembly Toolkit. See *Configuring the client for identity assertion authentication*: specifying the method if you have not previously specified this information.

**Configuring the server to handle identity assertion authentication**

Use this task to configure identity assertion authentication. The purpose of identity assertion is to assert the authenticated identity of the originating client from a Web service to a downstream Web service. Do not attempt to configure identity assertion from a pure client.

For the downstream Web service to accept the identity of the originating client (user name only), you must supply a special trusted BasicAuth credential that the downstream Web service trusts and can authenticate successfully. You must specify the user ID of the special BasicAuth credential in a trusted ID evaluator on the downstream Web service configuration. For more information on trusted ID evaluators, see *Trusted ID evaluators*. The server side passes the special BasicAuth credential into the trusted ID evaluator, which returns `true` or `false` that this ID is trusted. Once it is trusted, the user name of the client is mapped to the credential, which is used for authorization.

Complete the following steps to configure the server to handle identity assertion authentication information:

1. Launch the Assembly Toolkit.
2. Open the J2EE perspective by clicking **Window > Open Perspective > Other > J2EE**.
3. Select the Web services-enabled Enterprise JavaBean (EJB) or Web module.
4. In the Project Navigator window, locate the META-INF directory for an EJB module or the WEB-INF directory for a Web module.
5. Right-click the `webservices.xml` file, and click **Open With > Web Services Editor**.
6. Click the **Security Extensions** tab, which is located at the bottom of the Web services editor within the Assembly Toolkit.
7. Expand the **Request Receiver Service Configuration Details > Login Configuration** section. The options you can select are:
  - **BasicAuth**
  - **Signature**
  - **ID assertion**
  - **Lightweight Third Party Authentication (LTPA)**

8. Select **IDAssertion** to authenticate the client using the identity assertion data provided. The user ID of the client must be in the target user registry configured in WebSphere Application Server global security. You can select global security in the administrative console by clicking **Security > Global security**.

You can select multiple login configurations, which means that different types of security information can be received at the server. The order in which the login configurations are added determines the processing order when a request is received. Problems can occur if you have multiple login configurations added that have common security tokens. For example, ID assertion contains a BasicAuth token, which is the trusted token. For ID assertion to work properly, you must list ID assertion ahead of BasicAuth in the list or BasicAuth processing overrides ID assertion processing.

9. Expand the **IDAssertion** section and select both the **ID Type** and the **Trust Mode**.

- a. For ID Type, the options are:

- Username
- Distinguished name (DN)
- X509certificate

These choices are just preferences and are not guaranteed. Most of the time the Username option is used. You must choose the same ID Type as the client.

- b. For Trust Mode, the options are:

- BasicAuth
- Signature

The Trust Mode refers to the information sent by the client as the trusted ID.

- 1) If you select **BasicAuth**, the client sends basic authentication data (user ID and password). This basicauth data is authenticated to the configured user registry. When the authentication occurs successfully, the user ID must be part of the trusted ID evaluator trust list.
- 2) If you select **Signature**, the client signing certificate is sent. This certificate must be mappable to the configured user registry. For **Local OS**, the common name (CN) of the distinguished name (DN) is mapped to a user ID in the registry. For **Lightweight Directory Access Protocol (LDAP)**, the DN is mapped to the registry for the ExactDN mode. If it is in the CertificateFilter mode, attributes are mapped accordingly. In addition, the user name from the credential generated must be in the Trusted ID Evaluator trust list.

For more information on getting started with the Web Services Editor within the Assembly Toolkit , see [Configuring the server security bindings using the Assembly Toolkit](#).

After you specify how the server handles identity assertion authentication information, you must specify how the server validates the authentication information. See [Configuring the server to validate identity assertion authentication information](#) for more information.

## Configuring the server to validate identity assertion authentication information

Use this task to configure identity assertion authentication. The purpose of identity assertion is to assert the authenticated identity of the originating client from a Web service to a downstream Web service. Do not attempt to configure identity assertion from a pure client.

For the downstream Web service to accept the identity of the originating client (user name only), you must supply a special trusted BasicAuth credential that the downstream Web service trusts and can authenticate successfully. You must specify the user ID of the special BasicAuth credential in a trusted ID evaluator on the downstream Web service configuration. For more information on trusted ID evaluators, see [Trusted ID evaluators](#). The server side passes the special BasicAuth credential into the trusted ID evaluator, which returns a true or false response that this ID is trusted. Once it is trusted, the user name of the client is mapped to the credential, which is used for authorization.



Complete the following steps to validate the identity assertion authentication information:

1. Launch the Assembly Toolkit.
2. Open the J2EE perspective by clicking **Window > Open Perspective > Other > J2EE**.
3. Select the Web services-enabled Enterprise JavaBean (EJB) or Web module.
4. In the Project Navigator window, locate the META-INF directory for an EJB module or the WEB-INF directory for a Web module.
5. Right-click the webservices.xml file, and click **Open With > Web Services Editor**.
6. Click the **Binding Configurations** tab, which is located at the bottom of the Web services editor within the Assembly Toolkit.
7. Expand the **Request Receiver Binding Configuration Details > Login Mapping** section.
8. Click **Edit** to view the login mapping information. Click **Add** to add new login mapping information. The login mapping dialog is displayed. Select or enter the following information:

**Authentication method**

Specifies the type of authentication that occurs. Select **IDAssertion** to use basic authentication.

**Configuration name**

Specifies the Java Authentication and Authorization Service (JAAS) login configuration name. For the IDAssertion authentication method, enter `system.wssecurity.IDAssertion` for the Java Authentication and Authorization Service (JAAS) login configuration name.

**Use token value type**

Determines if you want to specify a custom token type. For the default authentication method selections, you do not need to specify this option.

**Token value type URI and Token value type local name**

When you select ID assertion, you cannot edit the token value type URI and local name values. Specifies custom authentication types. For the ID assertion authentication method, leave these values blank.

**Callback Handler Factory Class name**

Creates a JAAS CallbackHandler implementation that understands the following callbacks:

- `javax.security.auth.callback.NameCallback`
- `javax.security.auth.callback.PasswordCallback`
- `com.ibm.wsspi.wssecurity.auth.callback.BinaryTokenCallback`
- `com.ibm.wsspi.wssecurity.auth.callback.XMLTokenReceiverCallback`
- `com.ibm.wsspi.wssecurity.auth.callback.PropertyCallback`

For any of the default authentication methods (BasicAuth, IDAssertion, and Signature), use the callback handler factory default implementation. Enter the following class name for any of the default Authentication methods including IDAssertion:

```
com.ibm.wsspi.wssecurity.auth.callback.WSCallbackHandlerFactoryImpl
```

This implementation creates the correct callback handler for the default implementations.

**Callback handler factory property name and Callback handler factory property value**

Specifies callback handler properties for custom callback handler factory implementations. The default callback handler factory implementation does not need any specified properties. For ID assertion, leave these values blank.

**Login mapping property name and Login mapping property value**

Specifies properties for a custom login mapping. For the default implementations including IDAssertion, leave these values blank.

9. Expand the **Trusted ID Evaluator** section.



10. Click **Edit** to see a dialog that displays all the trusted ID evaluator information. The following table describes the purpose of this information.

**Class name**

Refers to the implementation of the trusted ID evaluator that you want to use. Enter the default implementation as `com.ibm.wsspi.wssecurity.id.TrustedIDEvaluatorImpl`. If you want to implement your own trusted ID evaluator, you must implement the `com.ibm.wsspi.wssecurity.id.TrustedIDEvaluator` interface.

**Property name**

Represents the name of this configuration. Enter `BasicIDEvaluator`.

**Property value**

Defines the name and value pairs that can be used by the trusted ID evaluator implementation. For the default implementation, the trusted list is defined here. When a request comes in and the trusted ID is verified, the user ID, as it appears in the user registry, must be listed in this property. Specify the property as a name and value pair where the name is `trustedId_n`. *n* is an integer starting from 0 and the value is the user ID associated with that name. An example list with the trusted names include two properties.

For example: `trustedId_0 = user1, trustedId_1 = user2`. The previous example means that both `user1` and `user2` are trusted. `user1` and `user2` must be listed in the configured user registry.

11. Expand the **Trusted ID Evaluator Reference** section.
12. Click **Enable** to add a new entry. The text you enter for the **Trusted ID Evaluator Reference** must be the same as the name entered previously in the **Trusted ID Evaluator**. Make sure that the name matches exactly because the information is case sensitive. If an entry is already specified, you can change it by clicking **Edit**.

You must specify how the server handles the identity assertion authentication method. See [Configuring the server to handle identity assertion authentication](#) if you have not previously specified this information.

## Securing Web services using signature authentication

WebSphere Application Server provides several different methods to secure your Web services. Extensible Markup Language (XML) digital signature is one of these methods. You might secure your Web services using any of the following methods:

- XML digital signature
- XML encryption
- Basicauth authentication
- Identity assertion authentication
- Signature authentication
- Pluggable token

With the signature authentication method, the request sender generates a signature security token using a callback handler. The security token returned by the callback handler is inserted in the Simple Object Access Protocol (SOAP) message. The request receiver retrieves the Signature security token from the SOAP message and validates it using a Java Authentication and Authorization Service (JAAS) login module. To use signature authentication to secure Web services, complete the following tasks:

1. Secure the client for signature authentication.
  - a. Configure the client for signature authentication: specifying the method.
  - b. Configure the client for signature authentication: collecting the authentication information.
2. Secure the server for signature authentication.
  - a. Configure the server to handle signature authentication.

- b. Configure the server to validate signature authentication information.

After completing these steps, you have secured your Web services using signature authentication.

## Configuring the client for signature authentication: specifying the method

This task is used to configure signature authentication. A signature refers to the use of an X.509 certificate to login on the target server. For more information on signature authentication, see Signature authentication method.

Complete the following steps to specify signature as the authentication method:

1. Launch the Assembly Toolkit.
2. Click **Windows > Open Perspective > J2EE** to access the Assembly Toolkit perspective.
3. Select the Web services-enabled Enterprise JavaBean(EJB) or Web module.
4. In the Project Navigator window, locate the META-INF directory for an EJB module or the WEB-INF directory for a Web module.
5. Right-click the webservicessclient.xml file, and click **Open With > Web Services Client Editor**.
6. Click the **Security Extensions** tab, which is located at the bottom of the Web Services Client Editor within the Assembly Toolkit.
7. Expand the **Request Sender Configuration > Login Configuration** section. The following login configuration options are valid for a managed client and Web services acting as a client are:

### **BasicAuth**

Use this option for a managed client.

### **Signature**

Use this option for a managed client.

### **IDAssertion**

Use this option for Web services acting as a client.

8. Select **Signature** to authenticate the client using the certificate used to digitally sign the request.

For more information on getting started with the Web services client editor within the Assembly Toolkit , see Configuring the client security bindings using the Assembly Toolkit.

After you specify signature as the authentication method, you must specify how to collect the authentication information. See Configuring the client for signature authentication: Collecting the authentication information for more information.

### **Signature authentication method:**

When using the signature authentication method, the security token is generated with a <ds:Signature> and a <wsse:BinarySecurityToken> element.

On the request sender side, a callback handler is invoked to generate the security token. On the request receiver side, a Java Authentication and Authorization Service (JAAS) login module is used to validate the security token. These two operations, token generation and token validation, are described in the following sections.

### **Signature token generation**

The request sender generates a Signature security token using a callback handler. The security token returned by the callback handler is inserted in the SOAP message. The callback handler is specified in the <LoginBinding> element of the bindings file, `ibm-webservicessclient-bnd.xml`. WebSphere Application Server provides the following callback handler implementation that can be used with the Signature authentication method:

```
com.ibm.wsspi.wssecurity.auth.callback.NonPromptCallbackHandler
```

You can add your own callback handlers that implement the `javax.security.auth.callback.CallbackHandler` implementation.

### Security token validation

The request receiver retrieves the Signature security token from the Simple Object Access Protocol (SOAP) message and validates it using a JAAS login module. The `<ds:Signature>` and `<wsse:BinarySecurityToken>` elements in the security token are used to perform the validation. If the validation is successful, the login module returns a Java Authentication and Authorization Service (JAAS) Subject. This Subject then is set as the identity of the running thread. If the validation fails, the request is rejected with a SOAP fault exception.

The JAAS login configuration is specified in the `<LoginMapping>` element of the bindings file. Default bindings are specified in the `ws-security.xml` file. However, you can override these bindings using the application-specific `ibm-webservices-bnd.xmi` file. The configuration information consists of a `CallbackHandlerFactory` and a `ConfigName`. The `CallbackHandlerFactory` specifies the name of a class that is used for creating the JAAS `CallbackHandler` object. WebSphere Application Server provides the `com.ibm.wsspi.wssecurity.auth.callback.WSCallbackHandlerFactoryImp` `CallbackHandlerFactory` implementation. The `ConfigName` specifies a JAAS configuration name entry. WebSphere Application Server searches in the `security.xml` file for a matching configuration name entry. If a match is not found, it searches the `wsjaas.conf` file. WebSphere Application Server provides the `system.wssecurity.Signature` default configuration entry, which is suitable for the signature authentication method.

### Configuring the client for signature authentication: collecting the authentication information

This task is used to configure signature authentication. A signature refers to the use of an X.509 certificate to login on the target server. For more information on signature authentication, see `Signature authentication method`.

Complete the following steps to specify how the client collects the authentication information for signature authentication:

1. Launch the Assembly Toolkit.
2. Open the J2EE perspective by clicking **Window > Open Perspective > Other > J2EE**.
3. Select the Web services-enabled Enterprise JavaBean (EJB) or Web module.
4. In the Project Navigator window, locate the META-INF directory for an EJB module or the WEB-INF directory for a Web module.
5. Right-click the `webservicesclient.xml` file, and click **Open With > Web Services Client Editor**.
6. Click the **Port Binding** tab, which is located at the bottom of the Web services client editor within the Assembly Toolkit.
7. Expand the **Security Request Sender Binding Configuration > Signing Information** and click **Edit** to modify the signing key name and signing key locator. To create new signing information, click **Enable**. The certificate that is sent to log in at the server is the one configured in the Signing Information section. Review the section on Key locators to understand how the signing key name maps to a key within the key locator entry.

The following list describes the purpose of this information. Some of these definitions are based on the XML-Signature specification, which is located at the following Web address:  
<http://www.w3.org/TR/xmlsig-core>

#### Canonicalization method algorithm

Canonicalizes the `SignedInfo` element before it is digested as part of the signature operation.

#### Digest method algorithm

Represents the algorithm that is applied to the data after transforms are applied, if specified, to yield the `<DigestValue>` element. The signing of the `DigestValue` element binds the resource

content to the signer key. The algorithm selected for the client request sender configuration must match the algorithm selected in the server request receiver configuration.

#### **Signature method algorithm**

Represents the algorithm that is used to convert the canonicalized <SignedInfo> value into the <SignatureValue> value. The algorithm selected for the client request sender configuration must match the algorithm selected in the server request receiver configuration.

#### **Signing key name**

Represents the key entry associated with the signing key locator. The key entry refers to an alias of the key, which is used to sign the request.

#### **Signing key locator**

Represents a reference to a key locator implementation. For more information on configuring key locators, see Key locators.

8. Expand the **Security Request Sender Binding Configuration > Login Binding** section.
9. Click **Edit** to view the login binding information. Then, select or enter the following information:

#### **Authentication method**

Specifies the type of authentication that occurs. Select **Signature** to use signature authentication.

#### **Token value type URI and Token value type URI local name**

When you select **Signature**, you cannot edit the Token value type URI and Local name values. Specifies custom authentication types. For signature authentication, leave these fields blank.

#### **Callback handler**

Specifies the Java Authentication and Authorization Server (JAAS) callback handler implementation for collecting signature information. Enter the following callback handler for signature authentication:

```
com.ibm.wsspi.wssecurity.auth.callback.NonPromptCallbackHandler
```

This callback handler is used because the signature method does not require user interaction.

#### **Basic authentication User ID and Basic authentication Password**

Leave the BasicAuth fields blank when Signature authentication is used.

#### **Property name and property value**

This field enables you to enter properties and name and value pairs for use by custom callback handlers. For signature authentication, you do not need to enter any information.

**Other customization entries:** There is a basic authentication entry in the Port Qualified Name Binding Details section. This entry is used for HTTP transport authentication, which might be required if the router servlet is protected.

Information specified in the Web services security signature authentication section overrides the basic authentication information specified in the Port Qualified Name Binding Details section for authorizing the Web service.

To use the signature authentication method, you must specify the authentication method in the **Login Config** section of the Assembly Toolkit. See Configuring the client for signature authentication: specifying the method if you have not previously specified this information.

## **Configuring the server to support signature authentication**

Use this task to configure signature authentication at the server. Signature authentication refers to an X.509 certificate sent by the client to the server. The certificate is used to authenticate to the user registry configured at the server. After a request is received by the server that contains the certificate, the server needs to log in to form a credential. The credential is used for authorization. If the certificate supplied

cannot be mapped to an entry in the user registry, an exception is thrown and the request ends without invoking the resource. For more information on signature authentication, see [Signature authentication method](#).

1. Launch the Assembly Toolkit.
2. Open the J2EE perspective by clicking **Window > Open Perspective > Other > J2EE**.
3. Select the Web services-enabled Enterprise JavaBean (EJB) or Web module.
4. In the Project Navigator window, locate the META-INF directory for an EJB module or the WEB-INF directory for a Web module.
5. Right-click the webservices.xml file, and click **Open With > Web Services Editor**.
6. Click the **Security Extensions** tab, which is located at the bottom of the Web Services Editor within the Assembly Toolkit.
7. Expand the **Request Receiver Service Configuration Details > Login Configuration** section. You can select from the following options:
  - BasicAuth
  - Signature
  - ID assertion
  - Lightweight Third Party Authentication (LTPA)
8. Select **Signature** to authenticate the client using an X509 certificate. The certificate that is sent from the client is the certificate that issued for signing the message. You must be able to map this certificate to the configured user registry. For Local operating system (OS) registries, the common name (cn) of the distinguished name (DN) is mapped to a user ID in the registry. For Lightweight Directory Access Protocol (LDAP), you can configure multiple mapping modes:
  - EXACT\_DN is the default mode that directly maps the DN of the certificate to an entry in the LDAP server.
  - CERTIFICATE\_FILTER is the mode that provides the LDAP advanced configuration with a place to specify a filter that maps specific attributes of the certificate to specific attributes of the LDAP server.

For more information on getting started with the Web services editor within the Assembly Toolkit, see [Configuring the server security bindings using the Assembly Toolkit](#).

After you specify how the server handles signature authentication information, you must specify how the server validates the authentication information. See [Configuring the server to validate signature authentication information](#) for more information.

## Configuring the server to validate signature authentication information

Use this task to configure signature authentication at the server. Signature authentication refers to an X.509 certificate sent by the client to the server. The certificate is used to authenticate to the user registry configured at the server. Once a request is received by the server that contains the certificate, the server needs to log in to form a credential. The credential is used for authorization. If the certificate supplied cannot be mapped to an entry in the user registry, an exception is thrown and the request ends without invoking the resource. For more information on signature authentication, see [Signature authentication method](#).

Complete the following steps to configure the server to validate signature authentication:

1. Launch the Assembly Toolkit.
2. Open the J2EE perspective by clicking **Window > Open Perspective > Other > J2EE**.
3. Select the Web services-enabled Enterprise JavaBean (EJB) or Web module.
4. In the Project Navigator window, locate the META-INF directory for an EJB module or the WEB-INF directory for a Web module.
5. Right-click the webservices.xml file, and click **Open With > Web Services Editor**.

6. Click the **Binding Configurations** tab, which is located at the bottom of the Web services editor within the Assembly Toolkit.
7. Expand the **Request Receiver Binding Configuration Details > Login Mapping** section.
8. Click **Edit** to view the login mapping information or click **Add** to add new login mapping information. The login mapping dialog is displayed and you select (or enter) the following information:

**Authentication method**

Specifies the type of authentication. Select **Signature** to use signature authentication.

**Configuration name**

Specifies the Java Authentication and Authorization Service (JAAS) login configuration name. For the signature authentication method, enter `system.wssecurity.Signature` for the JAAS login configuration name. This specification logs in with the `com.ibm.wsspi.wssecurity.auth.module.SignatureLoginModule` JAAS login module.

**Use token value type**

Determines if you want to specify a custom token type. For the default authentication method selections, you can leave this field blank.

**URI and local name**

When you select Signature method, you cannot edit the token value type URI and local name values. Specifies custom authentication types. For signature authentication, you can leave this field blank.

**Callback handler factory class name**

Creates a JAAS CallbackHandler implementation that understands the following callback handlers:

- `javax.security.auth.callback.NameCallback`
- `javax.security.auth.callback.PasswordCallback`
- `com.ibm.wsspi.wssecurity.auth.callback.BinaryTokenCallback`
- `com.ibm.wsspi.wssecurity.auth.callback.XMLTokenReceiverCallback`
- `com.ibm.wsspi.wssecurity.auth.callback.PropertyCallback`

For any of the default authentication methods (BasicAuth, IDAssertion, and Signature), use the callback handler factory default implementation. Enter the following class name for any of the default authentication methods including signature:

```
com.ibm.wsspi.wssecurity.auth.callback.WSCallbackHandlerFactoryImpl
```

This implementation creates the correct callback handler for the default implementations.

**Callback handler factory property name and callback handler factory property value**

Specifies callback handler properties for custom callback handler factory implementations. You do not need to specify any properties for the default callback handler factory implementation. For signature, you can leave this field blank.

**Login mapping property name and login mapping property value**

Specifies properties for a custom login mapping to use. For the default implementations including signature, you can leave this field blank.

Specify how the server handles the signature authentication method. See *Configuring the server to handle signature authentication* if you have not previously specified this information.

## Token type overview

The proposed Web services security draft defined two types of security tokens:

- Username token
- Binary security token



A user name token consists of a user name and, optionally, password information. You can include a user name token directly in the <Security> header within the message. Binary tokens, such as X.509 certificates, Kerberos tickets, Lightweight Third Party Authentication (LTPA) tokens, or other non-XML formats, require a special encoding for inclusion. The Web services security specification describes how to encode binary security tokens such as X.509 certificates and Kerberos tickets, and it also describes how to include opaque encrypted keys. The specification also includes extensibility mechanisms that you can use to further describe the characteristics of the credentials that are included with a message.

WebSphere Application Server, Version 5.0.2 supports user name tokens, which include both user name and password for basic authentication and user name, which is used for identity assertion. The WebSphere Application Server, Version 5.0.2 binary security token implementation supports both X.509 certificates and LTPA binary security. You extend the implementation to generate other types of tokens. However, Kerberos tickets are not supported in WebSphere Application Server, Version 5.0.2. Each type of token is processed by a corresponding token generation and validation module. The binary token generation and validation modules are pluggable that is based on the Java Authentication and Authorization Service (JAAS) framework. For example, an arbitrary XML-based token format is supported using the JAAS pluggable framework. WebSphere Application Server, Version 5.0.2 does not support an XML-based token that is used in the SecurityTokenReference.

You can define the types of tokens that the message can accept in the deployment descriptor extension file, `ibm.webservices-ext.xmi`. A message receiver might support one or more types of security tokens. The following example shows that the receiver supports four types of security tokens:

**Important:** In the following code sample, several lines were split into multiple lines due to the width of the printed page. See the close bracket for an indication of where each line of code ends.

```
?xml version="1.0" encoding="UTF-8"?>
<com.ibm.etools.webservice.wsext:WsExtension xmi:version="2.0"
xmlns:xmi="http://www.omg.org/XMI"
xmlns:com.ibm.etools.webservice.wsext=
"http://www.ibm.com/websphere/appserver/schemas/5.0.2/wsext.xmi"
xmi:id="WsExtension_1052760331306" routerModuleName="StockQuote.war">
 <wsDescExt xmi:id="WsDescExt_1052760331306" wsDescNameLink="StockQuoteFetcher">
 <pcBinding xmi:id="PcBinding_1052760331326" pcNameLink="urn:xmltoday-delayed-quotes"
scope="Session">
 <serverServiceConfig
xmi:id="ServerServiceConfig_1052760331326" actorURI="myActorURI">
 <securityRequestReceiverServiceConfig
xmi:id="SecurityRequestReceiverServiceConfig_1052760331326">
 <loginConfig xmi:id="LoginConfig_1052760331326">
 <authMethods xmi:id="AuthMethod_1052760331326" text="BasicAuth"/>
 <authMethods xmi:id="AuthMethod_1052760331327" text="IDAssertion"/>
 <authMethods xmi:id="AuthMethod_1052760331336" text="Signature"/>
 <authMethods xmi:id="AuthMethod_1052760331337" text="LTPA"/>
 </loginConfig>
 </idAssertion xmi:id="IDAssertion_1052760331336" idType="Username" trustMode="Signature"/>
 </serverServiceConfig>
 </pcBinding>
 </wsDescExt>
</com.ibm.etools.webservice.wsext:WsExtension>
```

The message sender might choose one of the token types that are supported by the receiver when sending a message. You can define the type of token to be used by the sending side in the client descriptor extension file, `ibm-webservicesclient-ext.xmi`. The following example shows that the sender chooses to send a UsernameToken to the receiver:

**Important:** In the following code sample, several lines were split into multiple lines due to the width of the printed page. See the close bracket for an indication of where each line of code ends.

```
<?xml version="1.0" encoding="UTF-8"?>
<com.ibm.etools.webservice.wsext:WsClientExtension xmi:version="2.0"
xmlns:xmi="http://www.omg.org/XMI"
xmlns:com.ibm.etools.webservice.wsext=
"http://www.ibm.com/websphere/appserver/schemas/5.0.2/wsext.xmi"
xmi:id="WsClientExtension_1052760331496">
 <ServiceRefs xmi:id="ServiceRef_1052760331506" serviceRefLink="service/StockQuoteService">
 </ServiceRefs>
</com.ibm.etools.webservice.wsext:WsClientExtension>
```



```

 <portQnameBindings xmi:id="PortQnameBinding_1052760331506"
portQnameLocalNameLink="StockQuote">
 <clientServiceConfig xmi:id="ClientServiceConfig_1052760331506"
actorURI="myActorURI">
 <securityRequestSenderServiceConfig
xmi:id="SecurityRequestSenderServiceConfig_1052760331506" actor="myActorURI">
 <loginConfig xmi:id="LoginConfig_1052760331506" authMethod="BasicAuth"/>

```

## Username token element

You can use the UsernameToken element to propagate a user name and, optionally, password information. Also, you can use this token type to carry basic authentication information. Both a user name and a password are used to authenticate the message. A UsernameToken containing the user name is used in identity assertion, which establishes the identity of the user based on the trust relationship.

The following example shows the syntax of the UsernameToken element:

```

<UsernameToken Id="...">
 <Username>...</Username>
 <Password Type="...">...</Password>
</UsernameToken>

```

The Web services security specification defines the following password types:

### **wsse:PasswordText (default)**

This type is the actual password for the user name.

### **wsse:PasswordDigest**

The type is the digest of the password for the user name. The value is a base64-encoded SHA1 hash value of the UTF8-encoded password.

WebSphere Application Server supports the default PasswordText type. However, it does not support password digest because most user registry security policies do not expose the password to the application software.

The following example illustrates the use of the <UsernameToken> element:

```

<S:Envelope xmlns:S="http://www.w3.org/2001/12/soap-envelope"
 xmlns:wsse="http://schemas.xmlsoap.org/ws/2002/04/secext">
 <S:Header>
 ...
 <wsse:Security>
 <wsse:UsernameToken>
 <wsse:Username>Joe</wsse:Username>
 <wsse:Password>ILoveJava</wsse:Password>
 </wsse:UsernameToken>
 </wsse:Security>
 </S:Header>
</S:Envelope>

```

### **Nonce, a randomly generated token:**

*Nonce* is a randomly generated, cryptographic token used to prevent the theft of user name tokens used with SOAP messages. Nonce is used with the basicauth authentication method.

Without nonce, when a user name token is passed from one machine to another machine using a non-secure transport, such as HTTP, the token might be intercepted and used in a replay attack. The same key might be reused when the username token is transmitted between the client and the server, which leaves it vulnerable to attack. The user name token can be stolen even if you use XML digital signature and XML encryption.

To help eliminate these replay attacks, the <wsse:Nonce> and <wsu:Created> elements are generated within the <wsse:usernameToken> element and used to validate the message. The request receiver or response

receiver checks the freshness of the message to verify the difference between when the message is created and the current time falls within a specified time period. Also, WebSphere Application Server verifies that the receiver has not processed the token within the specified time period. These two features are used to lessen the chance that a user name token is used for a replay attack.

### **Configuring nonce for the application level:**

*Nonce* is a randomly generated, cryptographic token used to thwart the highjacking of username tokens used with SOAP messages. Nonce is used in conjunction with the BasicAuth authentication method.

This task provides instructions on how to configure nonce for the application level using the WebSphere Application Server administrative console. You can configure nonce at the application level and cell level. However, you must consider the order of precedence:

1. Application level
2. Server level

If you configure nonce on the application level and the server level, the values specified for the application level take precedence over the values specified for the server level.

1. Connect to the administrative console by typing `http://localhost:9090/admin` in your Web browser unless you have changed the port number.
2. Click **Servers > Application Servers > server1**.
3. Under Additional Properties, click **Web Services: Default Bindings for Web Services Security > Login Mappings > New**.
4. Specify (optional) a value, in seconds, for the **Nonce Maximum Age** field. This panel is optional and only valid if the BasicAuth authentication method is specified. If you specify another authentication method and attempt to specify values for this field, the following error message displays and you must remove the specified value:

Nonce is not supported for authentication methods other than BasicAuth.

If you specify BasicAuth, but do not specify values for the Nonce Maximum Age field, the Web services security run time searches for a Nonce Maximum Age value on the server level. If a value is not found on the server level, the run time searches the cell level. If a value is not found on either the server level or the cell level, the default is 300 seconds.

The value specified for the Nonce Maximum Age field indicates how long the nonce is valid. You must specify a minimum of 300 seconds, but the value cannot exceed the number of seconds specified for the Nonce Cache Timeout field for either the server level

You can specify the Nonce Cache Timeout value for the server level by completing the following steps:

- a. Click **Servers > Application Servers > server\_name**.
- b. Under Additional Properties, click **Web Services: Default bindings for Web Services Security**.

You can specify the Nonce Cache Timeout value for the cell level by clicking **Security > Web Services > Properties**.

5. Specify (optional) a value, in seconds, for the **Nonce Clock Skew** field. The value specified for the Nonce Clock Skew field specifies the amount of time, in seconds, to consider when the message receiver checks the timeliness of the value. This panel is optional and only valid if the BasicAuth authentication method is specified. If you specify another authentication method and attempt to specify values for this field, the following error message displays and you must remove the specified value:

Nonce is not supported for authentication methods other than BasicAuth.

If you specify BasicAuth, but do not specify values for the **Nonce Clock Skew** field, the Web services security run time searches for a Nonce Clock Skew value on the server level. If a value is not found on the server level, the run time searches the cell level. If a value is not found on either the server level or the cell level, the default is 0 seconds.

Consider the following information when you set this value:

- Difference in time between the message sender and the message receiver if the clocks are not synchronized.
- Time needed to encrypt and transmit the message.
- Time needed to get through network congestion.

6. Restart the server.

### **Configuring nonce for the server level:**

*Nonce* is a randomly generated, cryptographic token used to prevent the theft of username tokens used with SOAP messages. Nonce is used in conjunction with the BasicAuth authentication method.

This task provides instructions on how to configure nonce for the server level using the WebSphere Application Server administrative console. You can configure nonce at the application level and the cell level. However, you must consider the order of precedence:

1. Application level
2. Server level

If you configure nonce on the application level and the server level, the values specified for the application level take precedence over the values specified for the server level. In a WebSphere Application Server environment, you must specify values for the Nonce Cache Timeout, Nonce Maximum Age, and Nonce Clock Skew fields on the server level to use nonce effectively. Complete the following steps to configure nonce on the server level:

1. Connect to the administrative console by typing `http://localhost:9090/admin` in your Web browser unless you have changed the port number.
2. Click **Servers > Application Servers > server1**.
3. Under Additional Properties, click **Web Services: Default Bindings for Web Services Security**.
4. Specify a value, in seconds, for the **Nonce Cache Timeout** field. The value specified for the Nonce Cache Timeout field indicates how long the nonce remains cached before it is expunged. You must specify a minimum of 300 seconds. However, if you do not specify a value, the default is 600 seconds. This field is required for the server level.
5. Specify (optional) a value, in seconds, for the **Nonce Maximum Age** field. The value specified for the Nonce Maximum Age field indicates how long the nonce is valid. You must specify a minimum of 300 seconds, but the value cannot exceed the number of seconds specified for the Nonce Cache Timeout field on the server level.

You can specify the Nonce Cache Timeout value for the cell level by clicking **Security > Web Services > Properties**.

This field is required for the server level.

6. Specify a value, in seconds, for the **Nonce Clock Skew** field. The value specified for the Nonce Clock Skew field specifies the amount of time, in seconds, to consider when the message receiver checks the timeliness of the value. Consider the following information when you set this value:
  - Difference in time between the message sender and the message receiver if the clocks are not synchronized.
  - Time needed to encrypt and transmit the message.
  - Time needed to get through network congestion.

You must specify at least 0 seconds for the Nonce Clock Skew field. However, the maximum value cannot exceed the number of seconds specified in the Nonce Maximum Age field on the server level. If you do not specify a value, the default is 0 seconds.

7. Restart the server. If you change the Nonce Cache Timeout value and do not restart the server, the change is not recognized by the server.

## Binary security token

The ValueType attribute identifies the type of the security token, for example, an LTPA token. The EncodingType indicates how the security token is encoded, for example, Base64Binary. The BinarySecurityToken element defines a security token that is binary encoded. The encoding is specified using the EncodingType attribute. The value type and space are specified using the ValueType attribute. The Web services security implementation for WebSphere Application Server, Version 5.0.2 supports both LTPA and X.509 certificate binary security tokens.

A binary security token has the following attributes that are used for interpretation:

- Value type
- Encoding type

The following example depicts an LTPA binary security token in a Web services security message header:

```
<wsse:BinarySecurityToken xmlns:ns7902342339871340177=
 "http://www.ibm.com/websphere/appserver/tokentype/5.0.2"
 EncodingType="wsse:Base64Binary"
 ValueType="ns7902342339871340177:LTPA">
 MIZ6LGPt2CzXBQfio9wZTo1VotWov0NW3Za6IU5K7Li78DSnIK6iHj3hxXgrUn6p4wZI
 8Xg26havepvmSJ8XxiACMihTJuh1t3ufsrjbFQJOqh5VcRvI+AKEaNmNtEgEV65jUYAC9
 C/iwBBWk5U/6DIk7LfXcTT0ZPAd+3D3nCS0f+6tnqMou8EG9mtMeTKccz/pJVTZjaRSo
 msu0sewsOKfl/WPsjW0bR/2g3NaVvBy18VITFBpUbGFVGgzHRjBKAGo+ctkl80nIVLlk
 TUjt/XdYvEpOr6QoddGi4okjDGPyoDxcvKZnReXww5UsoqlpfXwN4KG9as=
</wsse:BinarySecurityToken></wsse:Security></soapenv:Header>
```

As shown in the example, the token is Base64Binary encoded.

## XML token

XML tokens are offered in two formats, Security Assertion Markup Language (SAML) and Extensible rights Markup Language (XrML).

XML-based security tokens are growing in popularity. Two well-known formats are:

- Security Assertion Markup Language (SAML)
- Extensible rights Markup Language (XrML)

Using textensibility of the <wsse:Security> header in XML-based security tokens, you can directly insert these security tokens into the header.

SAML assertions are attached to Web services security messages using Web Services by placing assertion elements inside the <wsse:Security> header. The following example illustrates a Web services security message with a SAML assertion token.

```
<S:Envelope xmlns:S="...">&
 <wsse:Security xmlns:wsse="...">
 <saml:Assertion
 MajorVersion="1"
 MinorVersion="0"
 AssertionID="SecurityToken-ef375268"
 Issuer="elliottw1"
 IssueInstant="2002-07-23T11:32:05.6228146-07:00"
 xmlns:saml="urn:oasis:names:tc:SAML:1.0:assertion">
```

```

 ...
 </saml:Assertion>
 </wsse:Security>
</S:Header>
<S:Body>
 ...
</S:Body>
</S:Envelope>

```

For more information on SAML and XrML, see Resources for learning.

## Security token

A security token represents a set of claims made by a client that might include a name, password, identity, key, certificate, group, privilege, and so on.

Web services security provides a general-purpose mechanism to associate security tokens with messages for single message authentication. A specific type of security token is not required by Web services security. Web services security is designed to be extensible and support multiple security token formats to accommodate a variety of authentication mechanisms. For example, a client might provide proof of identity and proof of a particular business certification.

A security token is embedded in the SOAP message within the SOAP header. The security token within the SOAP header is propagated from the message sender to the intended message receiver. On the receiving side, the WebSphere Application Server security handler authenticates the security token and sets up the caller identity on the running thread.

## Securing Web services using a pluggable token

WebSphere Application Server provides several different methods to secure your Web services; a pluggable token is one of these methods. You might secure your Web services using any of the following methods:

- XML digital signature
- XML encryption
- Basicauth authentication
- Identity assertion authentication
- Signature authentication
- Pluggable token

Complete the following steps to secure your Web services using a pluggable token:

1. Generate a security token using the Java Authentication and Authorization Service (JAAS) CallbackHandler interface. The Web services security run time uses the JAAS CallbackHandler interface as a plug-in to generate security tokens on the client side or when Web services is acting as a client.
2. Configure your pluggable token. To use pluggable tokens to secure your Web services, you must configure both the client request sender and the server request receiver. You can configure your pluggable tokens using either the WebSphere Application Server administrative console or the WebSphere Application Server Toolkit. For more information, see the following topics:
  - Configuring pluggable tokens using the WebSphere Application Server Toolkit
  - Configuring pluggable tokens using the administrative console

## Configuring pluggable tokens using the Assembly Toolkit

This document describes how to configure a pluggable token in the request sender (`ibm-webservicesclient-ext.xmi` and `ibm-webservicesclient-bnd.xmi` file) and request receiver (`ibm-webservices-ext.xmi` and `ibm-webservices-bnd.xmi` file).

The pluggable token is required for the request sender and request receiver because they are a pair. The request sender and the request receiver must match for the receiver to accept a request.

Prior to completing these steps, it is assumed that you have already created a Web services-enabled Java 2 Platform, Enterprise Edition (J2EE) with a Web Services for J2EE (JSR 109) enterprise application. If not, see *Developing Web services to create Web services-enabled J2EE with a JSR 109 enterprise application*. See either of the following topics for an introduction of how to manage Web services security binding information for the server:

- Configuring the server security bindings using the Assembly Toolkit
- Configuring the server security bindings using the administrative console

You must specify the security constraints in the `ibm-webservicesclient-ext.xmi` and the `ibm-webservices-ext.xmi` files for the required tokens using the Assembly Toolkit.

Complete the following steps to configure the request sender using the `ibm-webservicesclient-ext.xmi` and `ibm-webservicesclient-bnd.xmi` files:

1. Launch the Assembly Toolkit.
2. Open the J2EE perspective by clicking **Window > Open Perspective > Other > J2EE**.
3. Select the Web services-enabled Enterprise JavaBean (EJB) or Web module.
4. In the Package Navigator window, locate the META-INF directory for an EJB module or the WEB-INF directory for a Web module.
5. Right-click the `webservicesclient.xml` file, and click **Open With > Web Services Client Editor**.
6. Click the **Security Extensions** tab. The Web service client security extensions editor is displayed.
  - a. Under Service References, select an existing service reference or click **Add** to create a new reference.
  - b. Under Port QName Bindings, select an existing port qualified name for the selected service reference or click **Add** to create a new port name binding.
  - c. Under Request Sender Configuration: Login Configuration, select an exiting authentication method or type in a new one in the editable list box (Lightweight Third Party Authorization (LTPA) is a supported token generation when Web services is acting as client).
  - d. Click **File > Save** to save the changes.
7. Click the **Web Services Client Binding** tab. The Web services client binding editor is displayed.
  - a. Under Port Qualified Name Binding, select an existing entry or click **Add** to add a new port name binding. The Web Services Client Binding editor displays for the selected port.
  - b. Under Login Binding, click **Edit** or **Enable**. The Login Binding dialog box is displayed.
    - 1) In the Authentication Method field, enter the authentication method. The authentication method that you enter in this field must match the authentication method defined on the Security Extension tab for the same Web service port. This field is mandatory.
    - 2) (Optional) Enter the token value type information in the **URI** and **Local name** fields. These fields are ignored for the BasicAuth, Signature, and IDAssertion authentication methods, but required for other authentication methods. The token value type information is inserted into the `<wsse:BinarySecurityToken>@ValueType` element for binary security token and is used as the namespace for the XML-based token.
    - 3) Enter an implementation of the Java Authentication and Authorization Service (JAAS) `javax.security.auth.callback.CallbackHandler` interface. This field is mandatory.



- 4) Enter the basic authentication information in the **User ID** and **Password** fields. The basic authentication information is passed to the construct of the CallbackHandler implementation. The use of the basic authentication information depends on the implementation of CallbackHandler.
  - 5) In the **Property** field, add name and value pairs. These pairs are passed to the construct of the CallbackHandler implementation as java.util.Map values.
  - 6) Click **OK**.  
Click **Disable** under Login Binding on the Web Services Client Port Binding tab to remove the authentication method login binding.
  - c. Click **File > Save** to save the changes.
8. In the Package Explorer window, right-click the webservices.xml file and click **Open With > Web Services Editor**. The Web Services window displays.
    - a. Click the **Security Extensions** tab. The Web Service Security Extensions editor is displayed.
      - 1) Under Web Service Description Extension, select an existing service reference or click **Add** to create a new extension.
      - 2) Under Port Component Binding, select an existing port qualified name for the selected service reference or click **Add** to create a new one.
      - 3) Under Request Receiver Service Configuration Details: Login Configuration, select an existing authentication method or click **Add** and enter a new method in the Add AuthMethod field that displays. You can select multiple authentication methods for the request receiver. The security token of the incoming message is authenticated against the authentication methods in the order that they are specified in the list. Click **Remove** to remove the selected authentication method or methods.
    - b. Click **File > Save** to save the changes.
    - c. Click the **Bindings** tab. The Web services bindings editor is displayed.
      - 1) Under Web Service Description Bindings, select an existing entry or click **Add** to add a new Web services descriptor.
      - 2) Click the **Binding Configurations** tab. The Web services binding configurations editor is displayed for the selected Web services descriptor.
      - 3) Under Request Receiver Binding Configuration Details: Login Mapping, click **Add** to create a new login mapping or click **Edit** to edit the selected login mapping. The Login mapping dialog is displayed.
        - a) In the Authentication method field, enter the authentication method. The information entered in this field must match the authentication method defined on the Security Extensions tab for the same Web service port. This field is mandatory.
        - b) In the Configuration name field, enter a JAAS login configuration name. You must define the JAAS login configuration name in the WebSphere Application Server administrative console under **Security > JAAS Configuration > Application Logins**. This is a mandatory field. For more information, see "Configuring Java Authentication and Authorization Service login" in the Information Center.
        - c) (Optional) Select **Use Token value type** and enter the token value type information in the URI and Local name fields. This information is optional for BasicAuth, Signature and IDAssertion authentication methods, but required for any other authentication method. The token value type is used to validate the <wsse:BinarySecurityToken>@ValueType element for binary security tokens and to validate the namespace of the XML-based token.
        - d) Under Callback Handler Factory, enter an implementation of the com.ibm.wsspi.wssecurity.auth.callback.CallbackHandlerFactory interface in the **Class name** field. This field is mandatory.
        - e) Under Callback Handler Factory property, click **Add** and enter the name and value pairs for the Callback Handler Factory Property. These name and value pairs are passed as



java.util.Map to the com.ibm.wsspi.wssecurity.auth.callback.CallbackHandlerFactory.init() method. The use of these name and value pairs is determined by the CallbackHandlerFactory implementation.

- f) Under Login Mapping Property, click **Add** and enter the name and value pairs for the Login mapping property. These name and value pairs are available to the JAAS Login Modules through the com.ibm.wsspi.wssecurity.auth.callback.PropertyCallback JAAS Callback interface. Click **Remove** to delete the selected login mapping.
  - g) Click **OK**.
- d. Click **File > Save** to save the changes.

The previous steps define how to configure the request sender to create security tokens in the Simple Object Access Protocol (SOAP) message and to configure the request receiver to validate the security tokens found in the incoming SOAP message. WebSphere Application Server supports pluggable security tokens.

You can use the authentication method defined in the login bindings and login mappings to generate security tokens in the request sender and validate security tokens in the request receiver.

After you configure pluggable tokens, you must configure both the client and the server to support pluggable tokens. See the following topics to configure the client and the server:

- Configuring the client for LTPA token authentication: Specifying LTPA token authentication
- Configuring the client for LTPA token authentication: Collecting the authentication information
- Configuring the server to handle LTPA token authentication
- Configuring the server to validate LTPA token authentication information

## Configuring pluggable tokens using the Administrative Console

Prior to completing these steps, it is assumed that you have already created a Web services-enabled Java 2 Platform, Enterprise Edition (J2EE) with a Web Services for J2EE (JSR 109) enterprise application. If not, see Developing Web services to create Web services-enabled J2EE with a JSR 109 enterprise application. See either of the following topics for an introduction of how to manage Web services security binding information for the server:

- Configuring the server security bindings using the Assembly Toolkit
- Configuring the server security bindings using the administrative console

This document describes how to configure a pluggable token in the request sender (ibm-webservicesclient-ext.xmi and ibm-webservicesclient-bnd.xmi file) and request receiver (ibm-webservices-ext.xmi and ibm-webservices-bnd.xmi file).

**Important:** The pluggable token is required for the request sender and request receiver as they are a pair. The request sender and the request receiver must match for a request to be accepted by the receiver.

Prior to completing these steps, it is assumed that you deployed a Web services-enabled enterprise application to the WebSphere Application Server.

Use the following steps to configure the client-side request sender (ibm-webservicesclient-bnd.xmi file) or server-side request receiver (ibm-webservices-bnd.xmi file) using the WebSphere Application Server Administrative Console.

1. Click **Applications > Enterprise Applications > enterprise\_application**.
2. Under Related Items, click either **EJB Modules > Uri** or **Web Modules > Uri**. The *Uri* is the Web services-enabled module

- a. Under Additional Properties, click **Web Services: Client Security Bindings** to edit the response sender binding information, if Web services is acting as client.
  - 1) Under Response Sender Binding, click **Edit**.
  - 2) Under Additional Properties, click **Login Binding**.
  - 3) Select **Dedicated Login Binding** to define a new login binding.
    - a) Enter the authentication method, this must match the authentication method defined in IBM extension deployment descriptor. The authentication method must be unique in the binding file.
    - b) Enter an implementation of the JAAS `javax.security.auth.callback.CallbackHandler` interface.
    - c) Enter the basic authentication information (User ID and Password) and the basic authentication information is passed to the construct of the `CallbackHandler` implementation. The usage of the basic authentication information is up to the implementation of the `CallbackHandler`.
    - d) Enter the token value type, it is optional for `BasicAuth`, `Signature` and `IDAssertion` authentication methods but required for any other authentication method. The token value type is inserted into the `<wsse:BinarySecurityToken>@ValueType` for binary security token and used as the namespace of the XML based token.
    - e) Click **Properties**. Define the property with name and value pairs. These pairs are passed to the construct of the `CallbackHandler` implementation as `java.util.Map`.  
Select **None** to deselect the login binding.
- b. Under Additional Properties, click **Web Services: Server Security Bindings** to edit the request receiver binding information.
  - 1) Under Request Receiver Binding, click **Edit**.
  - 2) Under Additional Properties, click **Login Mappings**.
  - 3) Click **New** to create new login mapping.
    - a) Enter the authentication method, this must match the authentication method defined in the IBM extension deployment descriptor. The authentication method must be unique in the login mapping collection of the binding file.
    - b) Enter a JAAS Login Configuration name. The JAAS Login Configuration must be defined in **Security > JAAS Configuration > Application Logins**. For more information, see "Configuring Java Authentication and Authorization Service login" in the Information Center.
    - c) Enter an implementation of the `com.ibm.wsspi.wssecurity.auth.callback.CallbackHandlerFactory` interface. This is a mandatory field.
    - d) Enter the token value type, it is optional for `BasicAuth`, `Signature` and `IDAssertion` authentication methods but required for any other authentication method. The token value type is used to validate against the `<wsse:BinarySecurityToken>@ValueType` for binary security token and against the namespace of the XML based token.
    - e) Enter the name and value pairs for the "Login Mapping Property" by clicking **Properties**. These name and value pairs are available to the JAAS Login Module or Modules by `com.ibm.wsspi.wssecurity.auth.callback.PropertyCallback` JAAS Callback. **Note:** This is true when editing existing login mappings but not when creating new login mappings.
    - f) Enter the name and value pairs for the "Callback Handler Factory Property", these name and value pairs is passed as `java.util.Map` to the `com.ibm.wsspi.wssecurity.auth.callback.CallbackHandlerFactory.init()` method. The usage of these name and value pairs is up to the `CallbackHandlerFactory` implementation.
  - c. Click authentication method link to edit the selected login mapping.
  - d. Click **Remove** to remove the selected login mapping or mappings.
3. Click **Save** in the upper-left section of the Administrative Console.

The previous steps define how to configure the request sender to create security tokens in the Simple Object Access Protocol (SOAP) message and the request receiver to validate the security tokens found in the incoming SOAP message. WebSphere Application Server supports pluggable security tokens.

You can use the authentication method defined in the login bindings and login mappings to generate security tokens in the request sender and validate security tokens in the request receiver.

Once you have configured pluggable tokens, you must configure both the client and the server to support pluggable tokens. See the following topics to configure the client and the server:

- Configuring the client for LTPA token authentication: specifying LTPA token authentication
- Configuring the client for LTPA token authentication: collecting the authentication information
- Configuring the server to handle LTPA token authentication
- Configuring the server to validate LTPA token authentication information

## Pluggable token support

You can extend the WebSphere Application Server login mapping mechanism to handle new types of authentication tokens. WebSphere Application Server provides a pluggable framework to generate security tokens on the sender-side of the message and to validate the security token on the receiver-side of the message. The framework is based on the Java Authentication and Authorization Service (JAAS) Application Programming Interfaces (APIs). Pluggable security token support provides plug-in points to support customer security token types including token generation, token validation, and client identity mapping to a WebSphere Application Server identity that is used by the Java 2 platform, Enterprise Edition (J2EE) authorization engine. Moreover, the pluggable token generation and validation framework supports XML-based tokens to be inserted into the Web service message header and validated on the receiver-side validation.

Use the `javax.security.auth.callback.CallbackHandler` implementation to create a new type of security token following these guidelines:

- Use a constructor that takes a user name (a string or null, if not defined), a password (a `char[]` or null, if not defined) and `java.util.Map` (empty, if properties are not defined).
- Use `handle()` methods that can process the following implementations:
  - `javax.security.auth.callback.NameCallback`
  - `javax.security.auth.callback.PasswordCallback`
  - `com.ibm.wsspi.wssecurity.auth.callback.XMLTokenCallback`
  - `com.ibm.websphere.security.auth.callback.WSCredTokenCallbackImpl`

If:

1. Either the `javax.security.auth.callback.NameCallback` or the `javax.security.auth.callback.PasswordCallback` implementation is populated with data, then a `<wsse:UsernameToken>` element is created.
2. `com.ibm.websphere.security.auth.callback.WSCredTokenCallbackImpl` is populated, the `<wsse:BinarySecurityToken>` element is created from the `com.ibm.websphere.security.auth.callback.WSCredTokenCallbackImpl` implementation.
3. `com.ibm.wsspi.wssecurity.auth.callback.XMLTokenCallback` is populated, a XML-based token is created based on the Document Object Model (DOM) element that is returned from the `XMLTokenCallback`.

Encode the token byte by using the security handler and not by using the `javax.security.auth.callback.CallbackHandler` implementation.

You can implement the `com.ibm.wsspi.wssecurity.auth.callback.CallbackHandlerFactory` interface, which is a factory for instantiating the `javax.security.auth.callback.CallbackHandler` implementation. For your own implementation, you must provide the `javax.security.auth.callback.CallbackHandler` interface. The Web service security run time instantiates the factory implementation class and passes the authentication

information from the Web services message header to the factory class through the setter methods. The Web services security run time then invokes the `newCallbackHandler()` method of the factory implementation class to obtain an instance of the `javax.security.auth.CallbackHandler` object. The object is passed to the JAAS login configuration.

The following is an example the definition of the `CallbackHandlerFactory` interface:

```
public interface com.ibm.wsspi.wssecurity.auth.callback.CallbackHandlerFactory {
 public void setUsername(String username);
 public void setRealm(String realm);
 public void setPassword(String password);
 public void setHashMap(Map properties);
 public void setTokenByte(byte[] token);
 public void setXMLToken(Element xmlToken);
 public CallbackHandler newCallbackHandler();
}
```

## Configuring the client for LTPA token authentication: specifying LTPA token authentication

Use this task to configure Lightweight Third-Party Authentication (LTPA) token authentication. Only configure the client for LTPA token authentication if the authentication mechanism configured in WebSphere Application Server is LTPA. When a client authenticates to a WebSphere Application Server, the credential created contains an LTPA token. When a Web service calls a downstream Web service, you can configure the first Web service to send the LTPA token from the originating client. Do not attempt to configure LTPA from a pure client. LTPA works only when you configure the client-side of a Web service acting as a client to a downstream Web service. For the downstream Web service to validate the LTPA token, the LTPA keys on both servers must be the same.

Complete the following steps to specify LTPA token as the authentication method:

1. Launch the Assembly Toolkit.
2. Open the J2EE perspective by clicking **Window > Open Perspective > Other > J2EE**.
3. Select the Web services-enabled Enterprise JavaBean (EJB) or Web module.
4. In the Project Navigator window, locate the META-INF directory for an EJB module or the WEB-INF directory for a Web module.
5. Right-click the `webservicesclient.xml` file, and click **Open With > Web Services Client Editor**.
6. Click the **Security Extensions** tab, which is located at the bottom of the Web services client editor within the Assembly Toolkit.
7. Expand the **Request Sender Configuration > Login Configuration** section.
8. Select **LTPA** as the authentication method. For more conceptual information on LTPA authentication, see LTPA.

After you specify LTPA token as the authentication method, you must specify how to collect the LTPA token information. See [Configuring the client for LTPA token authentication: Collecting the authentication information](#) for more information.

## Configuring the client for LTPA token authentication: Collecting the authentication method information

Use this task to configure Lightweight Third-Party Authentication (LTPA) token authentication. Do not configure the client for LTPA token authentication unless the authentication mechanism configured in WebSphere Application Server is LTPA. When a client authenticates to a WebSphere Application Server, the credential created contains an LTPA token. When a Web service calls a downstream Web service, you can configure the first Web service to send the LTPA token from the originating client. Do not attempt to configure LTPA from a pure client. LTPA works only when you configure the client-side of a Web service acting as a client to a downstream Web service. In order for the downstream Web service to validate the LTPA token, the LTPA keys on both servers must be the same.

Complete the following steps to specify how to collect the LTPA token authentication information:

1. Launch the Assembly Toolkit.
2. Open the J2EE perspective by clicking **Window > Open Perspective > Other > J2EE**.
3. Select the Web services-enabled Enterprise JavaBean (EJB) or Web module.
4. In the Project Navigator window, locate the META-INF directory for an EJB module or the WEB-INF directory for a Web module.
5. Right-click the webservicessclient.xml file, and click **Open With > Web Services Client Editor**.
6. Click the **Port Binding** tab, which is located at the bottom of the Web services client editor within the Assembly Toolkit.
7. Expand the **Security Request Sender Binding Configuration > Login Binding** section.
8. Click **Edit** to view the login binding information and select **LTPA**. If LTPA is not already there, enter it as an option. The login binding dialog is displayed. Select or enter the following information:

#### **Authentication method**

Specifies the type of authentication that occurs. Select **LTPA** to use identity assertion.

#### **Token value type URI and token value type local name**

When you select **LTPA**, you must edit the **token value type URI** and the **local name** fields. Specifies values for custom authentication types, which are authentication methods not mentioned in the specification. For the token value type URI field, enter the following string: `http://www.ibm.com/websphere/appserver/tokentype/5.0.2`. For the local name field, enter the following string: `LTPA`.

#### **Callback handler**

Specifies the Java Authentication and Authorization Service (JAAS) callback handler implementation for collecting the LTPA information. Specify the `com.ibm.wsspi.wssecurity.auth.callback.LTPATokenCallbackHandler` implementation for LTPA.

#### **Basic authentication user ID and basic authentication password**

For LTPA, you can leave these fields empty.

#### **Property name and property value**

For LTPA, you can leave these fields empty.

See Configuring the client for LTPA token authentication: Specifying LTPA token authentication if you have not previously specified this information.

## **Configuring the server to handle LTPA token authentication information**

This task is used to configure Lightweight Third-Party Authentication (LTPA). LTPA is a type of authentication mechanism in WebSphere Application Server security that defines a particular token format. The purpose of the LTPA token authentication is to flow the LTPA token from the first Web service, which authenticated the originating client, to the downstream Web service. Do not attempt to configure LTPA from a pure client. Once the downstream Web service receives the LTPA token, it validates the token to verify that the token has not been modified and has not expired. For validation to be successful, the LTPA keys used by both the sending and receiving servers must be the same.

Complete the following steps to specify that LTPA is authentication method. The authentication method indicated in these steps must match the authentication method specified for the client.

1. Launch the Assembly Toolkit.
2. Open the J2EE perspective by clicking **Window > Open Perspective > Other > J2EE**.
3. Select the Web services-enabled Enterprise JavaBean (EJB) or Web module.
4. In the Project Navigator window, locate the META-INF directory for an EJB module or the WEB-INF directory for a Web module.
5. Right-click the webservicess.xml file, select **Open With > Web Services Editor**.



6. Click the **Security Extensions** tab, which is located at the bottom of the Web services client editor within the Assembly Toolkit.
7. Expand the **Request Receiver Service Configuration Details > Login Configuration** section. You can select from the following options:
  - BasicAuth
  - Signature
  - ID assertion
  - LTPA
8. Select **LTPA** to authenticate the client using the LTPA token received from the request.

After you specify the authentication method, you must specify the information that the server must validate. See *Configuring the server to validate LTPA token authentication information* for more information.

## Configuring the server to validate LTPA token authentication information

This task is used to configure Lightweight Third-Party Authentication (LTPA). LTPA is a type of authentication mechanism in WebSphere Application Server security that defines a particular token format. The purpose of the LTPA token authentication is to flow the LTPA token from the first Web service, which authenticated the originating client, to the downstream Web service. Do not attempt to configure LTPA from a pure client. Once the downstream Web service receives the LTPA token, it validates the token to verify that the token has not been modified and has not expired. For validation to be successful, the LTPA keys used by both the sending and receiving servers must be the same.

Complete the following steps to specify how the server must validate the LTPA token authentication information:

1. Launch the Assembly Toolkit.
2. Open the J2EE perspective by clicking **Window > Open Perspective > Other > J2EE**.
3. Select the Web services-enabled Enterprise JavaBean (EJB) or Web module.
4. In the Project Navigator window, locate the META-INF directory for an EJB module or the WEB-INF directory for a Web module.
5. Right-click the webservices.xml file, and click **Open With > Web Services Editor**.
6. Click the **Binding Configurations** tab, which is located at the bottom of the Web services client editor within the Assembly Toolkit.
7. Expand the **Request Receiver Binding Configuration Details > Login Mapping** section.
8. Click **Edit** to view the login mapping information. The login mapping information is displayed. Select or enter the following information:

### Authentication method

Specifies the type of authentication that occurs. Select **LTPA** to use LTPA token authentication.

### Configuration name

Specifies the Java Authentication and Authorization Service (JAAS) login configuration name. For the LTPA authentication method, enter `WSLogin` for the JAAS login configuration name. This configuration understands how to validate an LTPA token.

### Use token value type

Determines if you want to specify a custom token type. For LTPA authentication, you must select this option because LTPA is considered a custom type. LTPA is not in the Web Services Security Specification.

### Token value type URI and local name

Specifies custom authentication types. If you select **Use Token value type** you must enter data into the Token value Type URI and local name fields. For the token value type URI field, enter the following string: `http://www.ibm.com/websphere/appserver/tokentype/5.0.2`. For the local name, enter the following string: `LTPA`

**Callback handler factory class name**

Creates a JAAS CallbackHandler implementation that understands the following callback handlers:

- `javax.security.auth.callback.NameCallback`
- `javax.security.auth.callback.PasswordCallback`
- `com.ibm.wsspi.wssecurity.auth.callback.BinaryTokenCallback`
- `com.ibm.wsspi.wssecurity.auth.callback.XMLTokenReceiverCallback`
- `com.ibm.wsspi.wssecurity.auth.callback.PropertyCallback`

For any of the default authentication methods (BasicAuth, IDAssertion, Signature, and LTPA), use the callback handler factory default implementation. Enter the following class name for any of the default authentication methods including LTPA:

```
com.ibm.wsspi.wssecurity.auth.callback.WSCallbackHandlerFactoryImpl
```

This implementation creates the correct callback handler for the default implementations.

**Callback handler factory property**

Specifies callback handler properties for custom callback handler factory implementations. Default callback handler factory implementation does not any property specifications. For LTPA, you can leave this field blank.

**Login mapping property**

Specifies properties for a custom login mapping. For default implementations including LTPA, you can leave this field blank.

See *Configuring the server to handle LTPA token authentication* if you have not previously specified this information.

***Lightweight Third Party Authentication:***

When you use the lightweight third party authentication (LTPA) method, the `<wsse:BinarySecurityToken>` security token generated is. On the request sender side, the security token is generated by invoking a callback handler. On the request receiver side, the security token is validated by a Java Authentication and Authorization Service (JAAS) login module.

The following information describes token generation and token validation operations.

**LTPA token generation**

The request sender uses a callback handler to generate an LTPA security token. The callback handler returns a security token that is inserted in the Simple Object Access Protocol (SOAP) message. Specify the appropriate callback handler in the `<LoginBinding>` element of the bindings file (`ibm-webservicesclient-bnd.xmi`). The following callback handler implementation can be used with the LTPA authentication method:

- `com.ibm.wsspi.wssecurity.auth.callback.LTPATokenCallbackHandler`

You can add your own callback handlers that implement the `javax.security.auth.callback.CallbackHandler` property.

When using the LTPA authentication method (or any authentication method other than BasicAuth, Signature or IDAssertion), the `TokenType` attribute of the `<LoginBinding>` element in the bindings file (`ibm-webservicesclient-bnd.xmi`) must be specified. The values to use for the LTPA `TokenType` attribute are:

- `uri="http://www.ibm.com/websphere/appserver/tokentype/5.0.2"`
- `localName="LTPA"`

**LTPA token validation**

The request receiver retrieves the LTPA security token from the SOAP message and validates the message using a JAAS login module. The `<wsse:BinarySecurityToken>` security token is used to



perform the validation. If the validation is successful, the login module returns a JAAS Subject. Subsequently, this Subject is set as the identity of the running thread. If the validation fails, the request is rejected with a SOAP fault.

The appropriate JAAS login configuration to use is specified in the bindings file <LoginMapping> element. Default bindings specified in the `ws-security.xml` file, but these can be overridden using the application-specific `ibm-webservices-bnd.xml` file. The configuration information consists of a `CallbackHandlerFactory`, a `ConfigName` and a `TokenValueType` attribute. The `CallbackHandlerFactory` specifies the name of a class to use to create the JAAS `CallbackHandler` object. A `CallbackHandlerFactory` implementation is provided (`com.ibm.wsspi.wssecurity.auth.callback.WSCallbackHandlerFactoryImpl`). The `ConfigName` attribute specifies a JAAS configuration name entry. The Web services security run time first searches the `security.xml` file for a matching entry and if a matching entry is not found, the run time searches the `wsjaas.conf` file. A default configuration entry suitable for the LTPA authentication method is provided (`WSLogin`). An appropriate `TokenValueType` element is located in the LTPA `LoginMapping` section of the default `ws-security.xml` file.

---

## Tuning Web services based on Web Services for J2EE

Performance considerations are the same for Web services applications that are developed and implemented based on the Web Services for Java 2 platform, Enterprise Edition (J2EE) and regular J2EE applications. See "Tuning performance" in the Information Center for more information about analyzing and tuning J2EE applications.

You can use the Performance Monitoring Infrastructure (PMI) to measure the time required to process Web services requests. To monitor Web services application performance:

1. Enable PMI services in application server through the administrative console. See "Enabling performance monitoring services in the application server through the administrative console" in the Information Center. Select the Web Service module, named `webServicesModule`, in step 7.
2. Monitor performance with Tivoli Performance Viewer. See "Monitoring performance with Tivoli Performance Viewer (formerly Resource Analyzer)" in the Information Center. In the left-hand pane of the performance view, expand the host and server and select **Web Services**. Run the Web services client application.

Measurements are available for the following items:

- Number of Web services loaded by the application server
- Number of requests received
- Number of requests dispatched to an implementation bean
- Number of requests dispatched with successful replies
- Average time in milliseconds between receiving the request and returning the reply
- Average time in milliseconds between receiving the request and dispatching it to the bean
- Average time in milliseconds between dispatch and receipt of reply from the bean
- Average time in milliseconds between receipt of reply from bean to return of result to client
- Average size of request and reply
- Average size of request
- Average size of reply

---

## Troubleshooting Web services based on Web Services for J2EE

Select the Web services topic area you want to troubleshoot:

- Command-line tools
- Java compiler errors
- Runtime errors and exceptions
- Client runtime errors and exceptions

- Serialization or deserialization errors

## Troubleshooting command-line tools for Web services based on Web Services for J2EE

This topic discusses troubleshooting command-line tools for Web services that are developed and implemented based on the Web Services for Java 2 platform, Enterprise Edition (J2EE) specification.

### WSDL2Java command-line tool

#### Using a proxy server to access the Internet while executing the WSDL2Java command causes your connection to time out

If you are using an environment that requires a proxy server to access the Internet during the execution of the **WSDL2Java** command, the **WSDL2Java** command might not find the Internet information because the proxy server has the potential to time out. For example, if the input `wsdl` file is located on the Internet instead of a local drive, and you need to retrieve it from the Internet, the **WSDL2Java** command fails to find it because the proxy server timed out.

You can work around this problem by editing the `WSDL2Java.bat` file when using a Windows operating system or the `WSDL2Java.sh` file if you are using a UNIX operating system. These files are located in the `<install_root>/WebSphere/AppServer/bin` directory.

If you are using a Windows operating system, set your proxy host and port values in the `WSDL2Java.bat` file:

```
PROXY_INFO="-Dproxy.httpHost=yourProxyHost -Dproxy.httpPort=yourProxyPort
```

If you are using a UNIX operating system, set your proxy host and port values in the `WSDL2Java.sh` file:

```
PROXY_INFO="-Dproxy.httpHost=yourProxyHost -Dproxy.httpPort=yourProxyPort
```

#### Emitter failure error occurs when running the WSDL2Java command on a WSDL document containing a JMS-style endpoint URL

If you run the **WSDL2Java** command-line tool on a WSDL document that contains a JMS-style endpoint URL, for example `jms:/...`, the `urlprotocols.jar` file that contains the custom protocol handler for the JMS protocol must be in the CLASSPATH. The error **WSWS3099E: Error: Emitter failure. Invalid endpoint address in port <x> in service <y>: <jms-url-string>** can be avoided by making sure the `urlprotocols.jar` file is in the CLASSPATH.

To add the `urlprotocols.jar` file to the CLASSPATH:

On Windows platforms, edit the `install_root\bin\setupCmdLine.bat` and locate the line which sets the `WAS_CLASSPATH` environment variable. Add `%install_root%\lib\urlprotocols.jar` to the end of the line that sets the `WAS_CLASSPATH` environment variable.

On UNIX platforms, edit the `install_root/bin/setupCmdLine.sh` file and add `$install_root/lib/urlprotocols.jar` to the end of the line that sets the `WAS_CLASSPATH` environment variable.

Make sure to use the proper delimiter character for your platform, for example, use a semi-colon (;) for Windows platforms and a colon (:) for UNIX platforms.

## Troubleshooting compiled bindings for Web services based on Web Services for J2EE

This topic discusses troubleshooting compiled bindings of Web services that are developed and implemented based on the Web Services for Java 2 platform, Enterprise Edition (J2EE) specification.

### Context root not recognized when mapping the default XML namespace to a Java package

When you map the default XML namespace to a Java package the context root is not recognized. If two namespaces are the same up to the first slash, they are mapped to the same Java package. For example, the XML namespaces `http://www.ibm.com/foo` and `http://www.ibm.com/bar` both map to the Java package `www.ibm.com`. Use the `-NStoPkg` option of the **Java2WSDL** command to specify the package for the fully qualified namespace.

### Java code to WSDL mapping cannot be reversed back to the original Java code

If you find that a WSDL file you created with the **Java2WSDL** command-line tool cannot be compiled when regenerated into Java code using the **WSDL2Java** command-line tool it is because the Java API for XML-based remote procedure call (JAX-RPC) mapping from Java code to WSDL is not reversible back to the original Java code.

To troubleshoot this problem review the WSDL file that was generated by the **Java2WSDL** tool using the information in Mapping between Java, WSDL and XML and the JAX-RPC specification available through Web services: Resources for learning. Use this information to determine which elements in the WSDL file are causing the problem. You can modify the WSDL file, or the original Java interface used to generate the WSDL file, and run the **Java2WSDL** command again.

## Troubleshooting the run time of Web services based on Web Services for J2EE

This topic discusses troubleshooting the run time of Web services that are developed and implemented based on the Web Services for Java 2 platform, Enterprise Edition (J2EE) specification.

You can troubleshoot run time errors and exceptions as follows:

- Trace SOAP messages
- Trace the components of Web services based Web Services for J2EE

### Tracing SOAP messages

This topic discusses tracing SOAP messages that request Web services that are developed and implemented based on the Web Services for Java 2 platform, Enterprise Edition (J2EE) specification.

You can trace the SOAP messages exchanged between a client and the server using the **TCPMon** command tool. The **TCPMon** command redirects messages from one port to another and records them. The WebSphere Application Server listens on port 9080. To trace messages sent to the application server, the **TCPMon** command is configured to listen on port 9088 and redirect them to 9080. The client is redirected to use port 9088 to access the Web service.

Redirecting an application client to a different port is most easily done by changing the SOAP address in the client's Web Services Description Language (WSDL) file to use port 9088 and then running the **wsdeploy** command-line tool on the client enterprise archive (EAR) file to regenerate the service implementation.

You should confirm that the server providing the Web service is running. The following task is performed on the machine providing the Web service.

To trace SOAP messages in Web services:

1. Set up a development and unmanaged client execution environment for Web services based on Web Services for J2EE
2. Run the **java -Djava.ext.dirs=%WAS\_EXT\_DIRS%** command. A window labeled TCPMonitor displays.
3. Configure the TCPMonitor to listen on port 9088 and forward messages to port 9080.
  - a. In the **Listen Port #** field, enter 9088.
  - b. Click **Listener**
  - c. In the **TargetHostname** field, enter localhost.
  - d. In the **Target Port #** field, enter 9080.
  - e. Click **Add**.
  - f. Click on the **Port 9088** tab that displays on the top of the page.

The messages exchanged between the client and server appear in the TCPMonitor Request and Response pane.

Save the message data and analyze it.

## Tracing Web services components based on Web Services for J2EE

The following are tasks in which you can enable trace for Web services that are developed and implemented based on the Web Services for Java 2 platform, Enterprise Edition (J2EE) specification.

1. Enable trace for a Web services unmanaged client.
  - a. Create a trace properties file by copying `%install_root%\WebSphere\AppServer\properties\TraceSettings.properties` file to the same directory as your client application Java archive (JAR) file.
  - b. Edit the properties file and change the value of `traceFileName` to output the trace data. For example, `traceFileName=c:\\temp\\myAppClient.trc`.
  - c. Edit the properties file to remove `com.ibm.ejs.ras.*=all=enabled` and add `com.ibm.ws.webservices.*=all=enabled`.
  - d. Add the option `-DtraceSettingsFile=<trace_properties_file>` to the Java command-line used to run the client, where `trace_properties_file` represents the name of the properties file created in steps 1-2. For example, **java -DtraceSettingsFile=TraceSettings.properties myApp.myAppMainClass**.
2. Enable trace for a Web services managed client.
  - a. Invoke the **launchClient** command-line tool with the following options:  
`-CCtrace=com.ibm.ws.webservices.*=all=enabled-CCtracefile=traceFileName` For example:  
**`%install_root%\bin\launchClient MyAppClient.ear-  
CCtrace=com.ibm.ws.webservices.*=all=enabled -CCtracefile=myAppClient.trc`**

See `launchClient` tool for more information.
3. Enable trace for a Web Services for J2EE server application.
  - a. Start WebSphere Application Server.
  - b. Open the administrative console.
    - a. Click **Servers >Application Servers > server**.
    - a. Click **Diagnostic Trace Service**.
    - a. In the **Trace Specification** field, delete the text `*=all=enabled` and add `com.ibm.ws.webservices.*=all=enabled`.
    - b. Click **Save** and **Apply**.

For more information see "Enabling trace" in the Information Center.

## Troubleshooting the run time for a Web services client based on Web Services for J2EE

This topic discusses troubleshooting Web services clients that are developed and implemented based on the Web Services for Java 2 platform, Enterprise Edition (J2EE) specification.

### Malformed URL exception displays when running a client that uses a JMS-style endpoint URL

If you are using the `launchClient` command to run a managed or unmanaged client that uses a JMS-style endpoint URL, the `urlprotocols.jar` file that contains the custom protocol handler for the JMS protocol must be in the CLASSPATH. The malformed URL exception can be avoided by making sure the `urlprotocols.jar` file is in the CLASSPATH.

To add the `urlprotocols.jar` file to the CLASSPATH:

On Windows platforms, edit the `install_root\bin\setupCmdLine.bat` and locate the line which sets the `WAS_CLASSPATH` environment variable. Add `%install_root%\lib\urlprotocols.jar` to the end of the line that sets the `WAS_CLASSPATH` environment variable.

On UNIX platforms, edit the `install_root/bin/setupCmdLine.sh` file and add `$install_root/lib/urlprotocols.jar` to the end of the line that sets the `WAS_CLASSPATH` environment variable.

## Troubleshooting serialization and deserialization in Web services based on Web Services for J2EE

The following are problems you might encounter performing serialization and deserialization in Web services that are developed and implemented based on the Web Services for Java 2 platform Enterprise Edition (J2EE) specification.

### Time zone information in deserialized `java.util.Calendar` is not as expected

When the client and server are based on Java code and a `java.util.Calendar` is received, the time zone in the received `java.util.Calendar` instance might be different from that of the `java.util.Calendar` instance that was sent.

This occurs because `java.util.Calendar` is encoded as an `xsd:dateTime` for transmission. An `xsd:dateTime` is required to encode the correct time (base time plus or minus a time zone offset), but is not required to preserve locale information, including the original time zone.

The fact that the time zone for the current locale is not preserved needs to be accounted for when comparing `Calendar` instances. The `java.util.Calendar` class `equals` method checks that the time zones are the same when determining equality. Since the time zone in a deserialized `Calendar` instance might not match the current locale, the `before` and `after` comparison methods should be used to test that two `Calendars` refer to the same date and time as shown below:

```
java.util.Calendar c1 = ...// Date and time in time zone 1
java.util.Calendar c2 = ...// Same date and equivalent time, but in time zone 2

// c1 and c2 are not equal because their time zones are different
if (c1.equals (c2)) System.out.println("c1 and c2 are equal");

// but c1 and c2 do compare as "not before and not after" since they represent
the same date and time
if (!c1.after(c2) & !c1.before(c2) {
 System.out.println("c1 and c2 are equivalent");
}
```

## Mixing Web services client and server bindings causes exceptions

Web Services for J2EE and Java API for XML-based remote procedure call (JAX-RPC) do not support "round-trip" mapping between Java code and a Web Services Description Language (WSDL) document for all Java types. For example, you cannot turn (serialize) a Java Date into XML code and then turn it back (deserialize) into a Java Date. It deserializes as Java Calendar.

If you have a Java implementation that you create a WSDL document from, and you generate client bindings from the WSDL document, the client classes can be different from the server classes even though the client classes have the same package and class names. The Web service client classes must be kept separate from the Web service server classes. For example, do not place the Web service server bindings classes in a utility Java archive (JAR) file and then include a Web service client JAR file that references the same utility JAR file.

If you do not keep the Web service client and server classes separate, a variety of exceptions can occur, depending on the Java classes used. The following is a sample stack trace error that can occur:

```
com.ibm.ws.webservices.engine.PivotHandlerWrapper TRAS0014I: The following exception was logged
 java.lang.NoSuchMethodError: com.ibm.wssvt.acme.websvcs.ExtWSPolicyData:
 method getStartDate()Ljava/util/Date; not found
at com.ibm.wssvt.acme.websvcs.ExtWSPolicyData_Ser.addElement(ExtWSPolicyData_Ser.java: 210)
at com.ibm.wssvt.acme.websvcs.ExtWSPolicyData_Ser.serialize (ExtWSPolicyData_Ser.java:29)
at com.ibm.ws.webservices.engine.encoding.SerializationContextImpl.serializeActual
 (SerializationContextImpl.java 719)
at com.ibm.ws.webservices.engine.encoding.SerializationContextImpl.serialize
 (SerializationContextImpl.java: 463)
```

The problem is caused by using an interface like the following for the Service Endpoint Interface in the service implementation:

```
package server;
public interface Test_SEI extends java.rmi.Remote {
 public java.util.Calendar getCalendar () throws java.rmi.RemoteException;
 public java.util.Date getDate() throws java.rmi.RemoteException;
}
```

When this interface is compiled and run through the **Java2WSDL** command-line tool, the WSDL document maps the methods as follows:

```
<wsdl:message name="getDateResponse">
 <wsdl:part name="getDateReturn" type="xsd:dateTime"/>
</wsdl:message>

<wsdl:message name="getCalendarResponse">
 <wsdl:part name="getCalendarReturn" type="xsd:dateTime"/>
</wsdl:message>
```

The JAX-RPC mapping implemented by the **Java2WSDL** tool has mapped both `java.util.Date` and `java.util.Calendar` to the XML type `xsd:dateTime`. The next step is to use the generated WSDL file to create a client for the Web service. When you run the **WSDL2Java** command-line tool on the generated WSDL, the generated classes include a different version of `server.Test_SEI`, for example:

```
package server;
public interface Test_SEI extends java.rmi.Remote {
 public java.util.Calendar getCalendar() throws java.rmi.RemoteException;
 public java.util.Date getDate() throws java.rmi.RemoteException;
}
```

**Note:** The client version of the `server.Test_SEI` interface is different from the server version in that both `getCalendar` and `getDate` methods return `java.util.Calendar`. The serialization and deserialization code that the client expects is the client version of the SEI. If the server version inadvertently appears in the client's CLASSPATH, at either compilation or execution time, an exception occurs.



In addition to the `NoSuchMethod` error, the `IncompatibleClassChangeError` and `ClassCastException` can occur, however, almost any run-time exception can occur. The best practice is to be diligent about separating client Web services bindings classes from server Web services bindings classes. The client bindings classes and server bindings classes should never be placed in the same module and, if they are in the same application, should not have bindings classes in utility JAR files that are shared between modules.

## Frequently asked questions about Web services based on Web Services for J2EE

This topic presents frequently asked questions about Web services that are developed and implemented based on the Web Services for Java 2 platform, Enterprise Edition (J2EE) specification.

- What IBM development tools work with Web Services for J2EE?
- Is Web Services for J2EE part of the J2EE specification?
- What is the relationship between Web Services for J2EE and the Web Service Invocation Framework (WSIF)?
- What is the relationship between Apache SOAP 2.3 and Web Services for J2EE?
- What is the relationship between the Apache Axis component of the Web services technology preview available with WebSphere Application Server 5.0 and Web Services for J2EE?
- What standards does the Web Services for J2EE component of WebSphere Application Server 5.0 support?
- Does Web Services for J2EE interoperate with other SOAP implementations, like .NET?
- Why can I not use a JavaBean to implement a SOAP Java Messaging Service (JMS) service?
- Does the SOAP JMS support interoperate with other vendors?
- How does two-way messaging with SOAP JMS work? Can it support multiple clients making simultaneous requests?

### What IBM development tools work with Web Services for J2EE?

WebSphere Studio Application Developer Version 5.1 and the Assembly Toolkit Version 5.1 both support the use of Web Services for J2EE. The Application Assembly Tool, included with WebSphere Application Server, and WebSphere Studio Application Developer versions earlier than Version 5.01, do not support Web Services for J2EE.

### Is Web Services for J2EE part of the J2EE specification?

For WebSphere Application Server 5.0.2, the Web Services for J2EE Version 1.0 specification is an addition to J2EE 1.3. J2EE 1.4 requires support for Web Services for J2EE Version 1.1. There are minor differences between the J2EE 1.3 Version (JSR-109 Version 1.0) and the J2EE 1.4 Version (JSR-109 Version 1.1).

### What is the relationship between Web Services for J2EE and the Web Service Invocation Framework (WSIF)?

Web Services for J2EE and WSIF represent two different programming models for accessing Web services. Web Services for J2EE is standard, Java-centric, and more statically bound to Web Services Description Language (WSDL) documents due to the use of generated stubs. WSIF directly models WSDL. WSIF is more suitable when dynamically interpreting WSDL. WebSphere Application Server Version 6.0 leverages both technologies to achieve dynamic, high performing standards-based Web services implementations.

### What is the relationship between Apache SOAP 2.3 and Web Services for J2EE?

Apache SOAP shipped with WebSphere Application Server Versions 4.0 and 5.0. It continues to co-exist with Web Services for J2EE. Apache SOAP is a proprietary API and applications written for it are not



portable to other SOAP implementations. Applications written for Web Services for J2EE should be portable to any vendor's implementation that supports Web Services for J2EE.

### **What is the relationship between the Apache Axis component of the Web services technology preview available with WebSphere Application Server 5.0 and Web Services for J2EE?**

The Web services technology preview leveraged the work that IBM contributed to the Apache Axis code base. The Web Services for J2EE support included with WebSphere Application Server 5.0.2 and Version 5.1 is derived from Apache Axis, but has diverged and contains many IBM-specific features to enhance performance, scalability, reliability, interoperability, and integration with the WebSphere Application Server.

### **What standards does the Web Services for J2EE component of WebSphere Application Server 5.0 support?**

The following standards are supported by the Web Services for J2EE component of WebSphere Application Server 5.0:

- SOAP Version 1.1
- Web Services Description Language (WSDL) Version 1.1
- Web Services for J2EE (JSR-109) Version 1.0
- Java API for XML-Based RPC (JAX-RPC) Version 1.0
- SOAP with attachments API for Java (SAAJ) Version 1.1

### **Does Web Services for J2EE interoperate with other SOAP implementations, like .NET?**

WebSphere Application Server Version 5.0.2 and Version 5.1 support Web services that are consistent with the the WS-I Basic Profile 1.0, and should interoperate with any other vendor conforming to this specification.

### **Why can I not use a Java bean to implement a SOAP Java Messaging Service (JMS) service?**

The SOAP JMS support uses Message Driven Beans (MDB) to implement the JMS endpoint. MDBs can only be used in the EJB container and delegate to an enterprise bean. If you want to use a Java bean instead of an enterprise bean to implement the service endpoint, you must create a "facade" enterprise bean that delegates to the Java bean.

### **Does the SOAP JMS support interoperate with other vendors?**

No. There is currently no specification for SOAP JMS, therefore each vendor chooses its own implementation technique.

### **How does two-way messaging with SOAP JMS work? Can it support multiple clients making simultaneous requests?**

Before a client issues a two-way request, it creates a temporary JMS queue to receive the response. This temporary queue is specified as the **replyTo** destination in the outgoing JMS request message. After the server processes the request, it directs the response to the **replyTo** destination specified in the request message. The client deletes the temporary queue after the response has been received. The server is able to handle simultaneous requests from multiple clients since each incoming request message contains the destination to which the reply should be sent.

---

## **Web services: Resources for learning**

This topic provides relevant supplemental information about the following Web services-related topics:

- Web services overview
  - Including the WebSphere Version 5 Web Services Handbook
- Developing Web services:

Including developing Web services based on the Java 2 platform, Enterprise Edition (J2EE) and Java API for XML-based remote procedure call (JAX-RPC) specifications.

- Web services gateway

Including an overview of the gateway and information about the WebSphere Application Server Edge components that are used for distributed gateway deployment.

- Universal Description Discovery and Integration (UDDI)

Including an overview about UDDI and information about the UDDI Java API.

- Web Services Invocation Framework (WSIF)

A look into the Apache Software Foundation and its maintenance of WSIF.

- SOAP

Including an overview about SOAP and information about the SOAP syntax and processing rules.

- Security

Including a roadmap to security, the WS-Security specification, best practices, a profile of the OASIS Security Assertion Markup Language (SAML) and more.

- Samples

Includes WebSphere Application Server Samples Gallery and Samples Central for Web services gateway, UDDI and WSIF.

- Other references

The information resides on IBM and non-IBM Internet sites, whose sponsors control the technical accuracy of the information.

These links are provided for convenience. Often, the information is not specific to the IBM WebSphere Application Server product, but is useful all or in part for understanding the product. When possible, links are provided to technical papers and Redbooks that supplement the broad coverage of the release documentation with in-depth examinations of particular product areas.

### **Web services overview**

- WebSphere Version 5 Web Services Handbook

This IBM Redbook describes the new concept of Web services from various perspectives. It presents the major building blocks Web services rely on. Well-defined standards and new concepts are presented and discussed.

- IBM Web Services architecture debuts

Introducing IBM Web services, a distributed software architecture of service components. This brief overview and in-depth interview on IBM DeveloperWorks cover the fundamental concepts of Web services architecture and what they mean for developers. The interview with IBM professional Rod Smith explores which types of developers Web services targets, how Web services reduces development time, what developers could be doing with Web services now, and takes a glance at the economics of dynamically discoverable services.

- Web services (r)evolution, Part 1

This article focuses on the benefits and challenges of building Web services applications. Web services might be an evolutionary step in designing distributed applications, however, they are not without their problems. Outlined are the difficulties developers face in creating a truly workable distributed system of Web services. This article also outlines author Graham Glass' plan for building peer-to-peer Web applications.

### **Developing Web services**

- JSR 109: Implementing Enterprise Web services

This document describes the Java 2 platform, Enterprise Edition (J2EE) specification.

- Java API for XML-based RPC (JAX-RPC): Core Web services API in the Java platform

This document reviews the JAX-RPC specification which enables Java technology developers to develop SOAP-based interoperable and portable Web services.

- **5.1+** A developer introduction to JAX-RPC, Part 1: Learn the ins and outs of the JAX-RPC type-mapping system. The Java API for XML-based Remote Procedure Calls (JAX-RPC) is an important step forward in the quest for Web services interoperability. This IBM developerWorks article explains the mapping between WSDL/XML types and Java types. It explains how the JAX-RPC standard defines this feature and some of the important points on designing an interoperable type system.
- **5.1+** A developer introduction to JAX-RPC, Part 2: Mine the JAX-RPC specification to improve Web service interoperability. This IBM developerWorks article explains how you can achieve the next level of Web service interoperability using the JAX-RPC standard client and server side interface definitions and message processing model. It includes information on developing JAX-RPC handlers and handler chains.
- **5.1+** Getting Started with JAX-RPC. This article explains some of the basic JAX-RPC programming concepts. It describes the JAX-RPC client and server programming models and provides some simple examples to illustrate their use. The article is intended to give developers a good grasp of how to use JAX-RPC to develop or use Web services.
- Web Services Description Language  
This article is a detailed overview of Web Services Description Language (WSDL), which includes programming specifications.

### Web services gateway

- The IBM Web services gateway: Technical Overview. A different version of the gateway is available as a component of a product called IBM WebSphere Business Connection. This brief technical summary from WebSphere Business Connection applies equally well to the version of the gateway in WebSphere Application Server.
- Information Center for WebSphere Application Server Edge components. This Information Center contains a library of PDF online books covering all aspects of the WebSphere Application Server Edge components. Scaling the Web services gateway builds upon the load-balancing capabilities of these components.

### UDDI

- Universal Description, Discovery and Integration  
This article is a detailed overview of Universal Description, Discovery and Integration (UDDI).
- UDDI4J: Matchmaking for Web services  
Reviewed in this article are the basics of UDDI, the Java API to UDDI, and how you can use this technology to start building, testing, and deploying your own Web services.

### WSIF

- The Apache Software Foundation. The Apache Software Foundation provides support for the Apache community of open-source software projects. Of particular interest is the Apache Web services project. The WSIF source code has been donated by IBM to the Apache Software Foundation, and is maintained here as an Apache project.

### SOAP

- SOAP  
This article is a detailed overview of SOAP, which includes programming specifications.
- SOAP Security Extensions: Digital Signature  
This document specifies the syntax and processing rules of a SOAP header entry to carry digital signature information within a SOAP 1.1 Envelope

### Security

- Security in a Web Services World: A Proposed Architecture and Roadmap

This document describes a proposed model for addressing security within a Web service environment. It defines a comprehensive Web Services Security model that supports, integrates, and unifies several popular security models, mechanisms, and technologies, including both symmetric and public key technologies, in a way that enables a variety of systems to securely interoperate in a platform and language-neutral manner. It also describes a set of specifications and scenarios that show how these specifications can be used together.

- Web Services Security (WS-Security)

The Web Services Security specifications describe enhancements to SOAP messaging to provide quality of protection through message integrity, message confidentiality, and single message authentication. These mechanisms can be used to accommodate a wide variety of security models and encryption technologies. Web Services Security also provides a general-purpose mechanism for associating security tokens with messages. Additionally, Web Services Security describes how to encode binary security tokens. Specifically, the specification describes how to encode X.509 certificates and Kerberos tickets, as well as how to include opaque encrypted keys. It also includes extensibility mechanisms that can be used to further describe the characteristics of the credentials that are included with a message.

- SOAP Security Extensions: Digital Signature

This document specifies the syntax and processing rules of a SOAP header entry to carry digital signature information within a SOAP 1.1 Envelope

- Web Services Security Addendum

This document describes clarifications, enhancements, best practices, and errata of the Web Services Security specification.

- WS-Security Profile of the OASIS Security Assertion Markup Language (SAML) Working Draft 04, 10 September 2002

This document proposes a set of standards for SOAP extensions used to increase message confidentiality.

- Web Services Security: SOAP Message Security Working Draft 12, Monday 21 April 2003

This document describes the support for multiple token formats, trust domains, signature formats, and encryption technologies.

- JSR 55:Certification Path API

This document provides a short description of the certification path API.

- XML-Signature Syntax and Processing

This document specifies XML digital signature processing rules and syntax. XML signatures provide integrity, message authentication, or signer authentication services for data of any type, whether located within the XML that includes the signature or elsewhere.

- Canonical XML Version 1.0

This specification describes a method for generating a physical representation, the canonical form, of an XML document that accounts for the permissible changes.

- Exclusive XML Canonicalization Version 1.0

Canonical XML [XML-C14N] specifies a standard serialization of XML that, when applied to a subdocument, includes the subdocument's ancestor context including all of the namespace declarations and attributes in the "xml:" namespace.

- XML Encryption Syntax and Processing

This document specifies a process for encrypting data and representing the result in XML.

- Decryption Transform for XML Signature

This document specifies an XML Signature "decryption transform" that enables XML Signature applications to distinguish between those XML Encryption structures that were encrypted before signing, and must not be decrypted, and those that were encrypted after signing, and must be decrypted, for the signature to validate.

- WS-Security

This document specifies resources for the April 2002 Web Services Security Specification. The following addenda and drafts are available:

– <http://schemas.xmlsoap.org/ws/2002/07/secext/>

- <http://schemas.xmlsoap.org/ws/2002/07/utility/>
- OASIS draft 12 for secext
- OASIS draft 12 for utility
- Specification: Web Services Security (WS-Security) Version 1.0 05 April 2002  
A major area of gateway security is based upon this emerging standard.
- XML Encryption Syntax and Processing W3C Recommendation 10 December 2002
- XML-Signature Syntax and Processing W3C Recommendation 12 February 2002
- Web Services Security Addendum
- Web Services Security Core Specification Working Draft 01, 20 September 2002
- Web Services Security: SOAP Message Security Working Draft 13, Thursday, 01 May 2003
- Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile, RFC3280, April 2002
- OASIS Web Services Security Technical Committee

### Samples

- Samples Gallery
- Samples Central. Samples and associated documentation for the following Web services components are available through the Samples Central page of the IBM WebSphere Developer Domain Web site:
  - The Web services gateway.
  - The IBM private UDDI registry.
  - The Web Services Invocation Framework (WSIF).

### Other references

- The Apache Software Foundation. The Apache Software Foundation provides support for the Apache community of open-source software projects. Of particular interest is the Apache Web services project.
- Web services insider, Part 1: Reflections on SOAP  
What is the current state of the *Web services revolution*? Find out at this Web site that features the column *Web services insider, Part 1*. The author answers this question by reviewing the tools and technologies that have emerged over the past year, highlighting their differences and similarities.
- The Web services insider, Part 2: A summary of the W3C Web Services Workshop  
This is a brief summary of a W3C Web services workshop.

---

## Chapter 9. Overview: Online garden retailer Web services scenarios

This set of scenarios is inspired by an online retailer called Plants by WebSphere. Plants by WebSphere uses Web services support in WebSphere Application Server to improve communications with its suppliers. The more advanced scenarios describe Web services support available only in particular editions of WebSphere Application Server. Consult your product documentation to confirm what is supported by your edition.

You might recognize Plants by WebSphere as a sample application available in the WebSphere Samples Gallery. These scenarios are loosely related. They describe how the fictional online retailer could use a variety of Web services technologies, some of which are beyond those currently demonstrated by the sample.

Web services are middleware. Using Web services you can connect applications together, no matter how each application is implemented or where it is located. For example, Web services can connect retailers to wholesale suppliers. Middleware is not new. What is new in Web services is that this connectivity is based upon open standards and Web technologies. Web services operate at a level of abstraction that is similar to the Internet, and they can work with any operating system, hardware platform or programming language that can be Web-enabled.

The Plants by WebSphere storefront sells plants and gardening supplies. As customers order merchandise, the site checks the merchandise availability in its inventory database. The scenarios show how the inventory system can grow in stages, using various Web services technologies to improve its capabilities.

- Before Web services

As featured in the Samples Gallery, the Plants by WebSphere application already has Web services capabilities. See below for a description of how the online garden retailer might have operated *prior* to adopting Web services technology. Key Web services components are introduced. To determine which components are available with the particular editions of WebSphere Application Server that you have purchased, consult the documentation for each edition.

- Static inquiry on supplier

In this scenario, the garden retailer turns the existing Web application into a Web service for checking the inventory of its main wholesale garden supplier.

- Dynamic inquiry on supplier

In this scenario, the garden retailer uses Web services to perform an inventory search of several wholesale suppliers.

- Cross supplier inquiry

In this scenario, the garden retailer makes its Web service available for use by others who need the service.

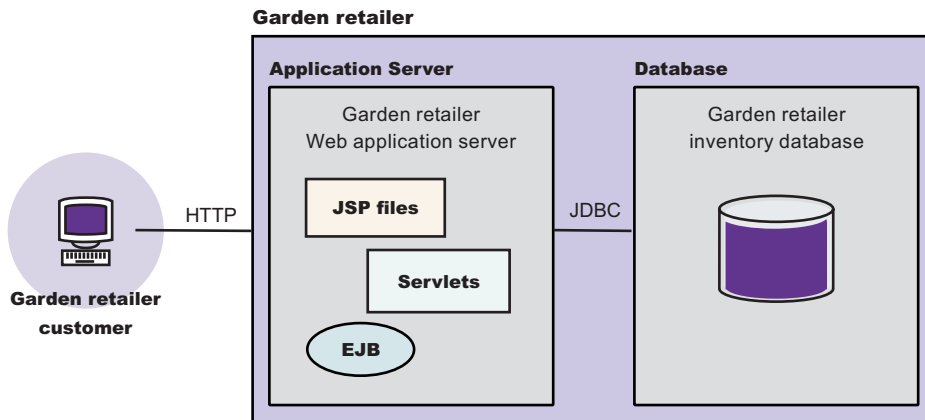
At present, these scenarios provide descriptions rather than step by step instructions. To gain experience with Web services coding, see the WebSphere Samples Gallery. It provides detailed instructions for building, configuring, and running the Plants by WebSphere sample application and others.

### Before Web services

Suppose that the Plants by WebSphere storefront does not use Web services. The garden retailer has established an impressive Internet storefront enabling customers to shop and order merchandise. To determine whether a customer order can be filled, Web applications rely on enterprise beans to query the Plants by WebSphere inventory database. If the item is in stock, the site confirms the order to the customer.



If a customer orders an item that is out of stock, the site notifies the customer that the item is out of stock, and encourages the customer to place the item on backorder. Later, long after the customer has left the Plants by WebSphere site, the site administrator or inventory manager might call or fax the supplier to obtain more inventory.



## Introducing Web services

Web services could give Plants by WebSphere an automated way to have out of stock items shipped to its warehouse or directly to customers. If suppliers can be contacted quickly enough, Plants by WebSphere does not have to inform its customers that the item was out of stock. Plants by WebSphere could begin to reduce its own inventory if doing so is a desirable business move.

Web services is built on the following core technologies:

- **XML (Extensible Markup Language)**

XML solves the problem of data independence. You use it to describe data, and also to map that data into and out of any application or programming language.

To have their applications exchange information such as merchandise price and availability, Plants by WebSphere and its supplier will put the data in a set of XML tags to which both parties agree.

For more information, see the XML specification on [www.w3.org](http://www.w3.org).

- **WSDL (Web Services Description Language)**

You use this XML-based language to create a description of an underlying application. It is this description that turns an application into a Web service, by acting as the interface between the underlying application and other Web-enabled applications.

Plants by WebSphere has an application capable of querying the supplier inventory. To enable communication with the supplier over the Internet, the company turns the application into a Web service.

For more information, see the WSDL specification on [www.w3.org](http://www.w3.org)

- **SOAP (Simple Object Access Protocol)**

SOAP is the core communications protocol for the Web, and most Web services use this protocol to talk to each other.

SOAP is an XML format for Web services requests. According to the SOAP specification, SOAP is "a lightweight protocol for exchange of information in a decentralized, distributed environment. It is an XML based protocol that consists of three parts: an envelope that defines a framework for describing what is in a message and how to process it, a set of encoding rules for expressing instances of application-defined datatypes, and a convention for representing remote procedure calls and responses. SOAP can potentially be used in combination with a variety of other protocols; however, the only bindings defined in this document describe how to use SOAP in combination with HTTP and HTTP Extension Framework.



Because they are external to the Plants by WebSphere intranet, communications with its suppliers will utilize SOAP over HTTP. Web services operating within the company intranet can use other transports, such as local Java bindings. The Web Services Invocation Framework (WSIF) component described below can help Plants by WebSphere applications dynamically choose the optimal transport mechanism for a given situation.

For more information, see the SOAP specification on [www.w3.org](http://www.w3.org).

- **Web Services for J2EE**

Web Services for J2EE, also known as JSR-109, defines how J2EE applications create and access Web services.

Chapter 8, “Using Web services based on Web Services for J2EE,” on page 395 describes how to implement a Web service interface to an existing application, then deploy your Web service within the application server.

- **Java API for Remote Procedure Calls**

JAX-RPC, also known as JSR-101, defines how Java applications access Web services.

The WebSphere product line provides these additional components to help you get the most out of your Web services. The scenarios describe in greater detail how Plants by WebSphere uses each one.

### **A private Universal Description, Discovery and Integration (UDDI) registry**

A private UDDI registry provides a way to publish and discover information about the Web services that are available within your organization. It does the same job for Web services that a business telephone directory does for business addresses and telephone numbers.

If you publish your Web service to UDDI, you make it available for other people to discover and reuse. Reuse of the service saves cost and effort, and publication minimizes the risk of duplicate services being developed.

You publish your Web service to UDDI after you have deployed it to the application server - unless you are also using the Web services gateway, in which case you use the gateway to publish the service to UDDI.

For more information, see Chapter 11, “Enabling the IBM WebSphere UDDI Registry,” on page 693.

### **A Web Services Invocation Framework (WSIF)**

SOAP bindings for Web services are part of the WSDL specification. So when you think of using a Web service, you probably think of assembling a SOAP message and sending it across the network to the service endpoint, using some SOAP client API. But if you invoke a Web service using WSIF, then a client application can dynamically choose the optimal transport mechanism to use to invoke Web service operations.

For example, a Web service provider might offer a SOAP binding for the service and a local Java binding that allows you to treat the local service implementation (a Java class) as a Web service. If the client is deployed in the same environment as the service, then the local Java binding for the service can be used. This provides more efficient communication with the service by making direct Java calls rather than using the SOAP binding.

To deploy a Web service as a WSIF-enabled service, you first develop and deploy the Web service, then you develop the WSIF client based on the WSDL document for that Web service - unless you are also using the Web services gateway, in which case the gateway automatically redeploys your Web service as a WSIF-enabled service.

For more information, see Chapter 10, “Web Services Invocation Framework (WSIF): Enabling Web services,” on page 651.

### **A Web services gateway**

You use the gateway to handle Web service invocations between Internet and Intranet environments. You use it to make your internal Web services available externally, and to make external Web services available to your internal systems. You also use it to specify:

- The transport mechanisms (or channels) on which messages can be carried to and from the service.
- The filters (if any) that act upon these incoming and outgoing messages.
- The UDDI registries (if any) to which you want the service to be published
- The levels of security that you want to apply to the service.

When you deploy a Web service to the gateway, the gateway creates a copy of the WSDL file for that service and stores it at a new Web address. Users of the service through the gateway then use the gateway copy of the WSDL file. So (if possible) you should decide whether or not you want to use the gateway before you make the Web addresses of your deployed services available to others.

For more information, see Chapter 12, “Web services gateway: Enabling Web services,” on page 823.

### **WS-Security (Web services security)**

Web services security for WebSphere Application Server is based on standards included in the Web services security (WS-Security) specification that address how to provide protection for messages exchanged in a Web service environment. It defines the core facilities for protecting the integrity and confidentiality of a message and provides mechanisms for associating security-related claims with the message. Web services security is a message-level standard based on securing SOAP messages through XML digital signature, confidentiality through XML encryption, and credential propagation through security tokens.

For more information, see “Securing Web services based on WS-Security” on page 483.

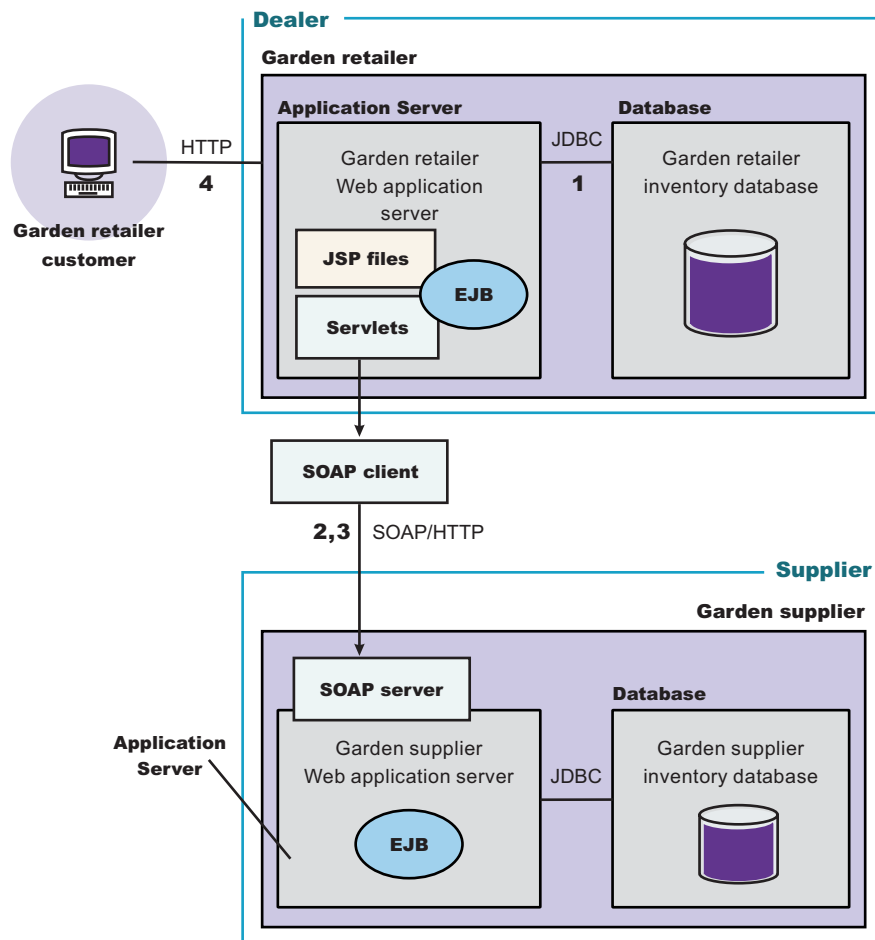
---

## **Web services scenario: Static inquiry on supplier**

In this scenario, an online supply retailer turns an application into a Web service for checking the inventory of its main supplier.

Plants by WebSphere is an Internet storefront that sells plants and gardening supplies. The company realizes that its existing Web application can be “wrapped” with Web services programming interfaces. For example, the BackOrderStock session bean can be turned into a SOAP Web services client that interacts with a Supplier Web services application located at the supplier. Specifically, the BackOrderStock session bean functionality is turned into a Web services for J2EE client and a JAX-RPC client. These two client

are J2EE and Java standards for Web services clients.



### How out of stock items are handled

The following events happen when a customer on the Plants by WebSphere site orders an item that is not available according to the Plants by WebSphere inventory:

1. Plants by WebSphere checks its own inventory.  
The application powering the Web site checks the Plants by WebSphere inventory database. It discovers that the item is not in stock.
2. Plants by WebSphere uses a Web service to check the supplier inventory.  
The application invokes a SOAP client that communicates with a SOAP server at the supplier site to ascertain whether the supplier has the item in stock. The supplier data is sent to Plants by WebSphere.
3. Plants by WebSphere either obtains the out of stock item, or does not.  
If the supplier indicates that the item is in stock, the application powering Plants by WebSphere determines whether to order the item on behalf of the customer. The exchange of data can include checking a price threshold above which Plants by WebSphere will not order the wholesale item. It could also include decision-making information such as how long the supplier will require to deliver the item, or a date on which the manufacturer plans to discontinue the item.
4. Plants by WebSphere notifies its customer of the outcome, as soon as possible.  
If the supplier can be consulted quickly enough, Plants by WebSphere does not have to bother its customer with concerns about availability. It simply confirms that the item is available, as though the

item is in stock at Plants by WebSphere. If the supplier inventory temporarily lacks the item, or Plants by WebSphere opted not to order the item from the supplier, Plants by WebSphere will issue an appropriate response to the customer.

See the Samples Gallery for details of the Supplier application, a business to business (B2B) application that uses a message-driven bean (MDB), a stateless session bean, and Java message service.

### **Web services technologies used in this scenario**

This scenario uses the following Web services technologies.

#### **XML (Extensible Markup Language)**

XML is used to standardize the exchange of data between Plants by WebSphere and its supplier.

#### **Web Services for J2EE**

Web Services for J2EE, also known as JSR-109, defines how J2EE applications create and access Web services.

#### **Java API for Remote Procedure Calls**

JAX-RPC, also known as JSR-101, defines how Java applications access Web services.

#### **WSDL (Web Services Description Language)**

WSDL is used to turn the existing application into a Web service, by acting as the interface between the underlying application and other Web-enabled applications.

#### **SOAP (Simple Object Access Protocol)**

SOAP is the protocol by which the Web service communicates with the supplier over the Internet.

---

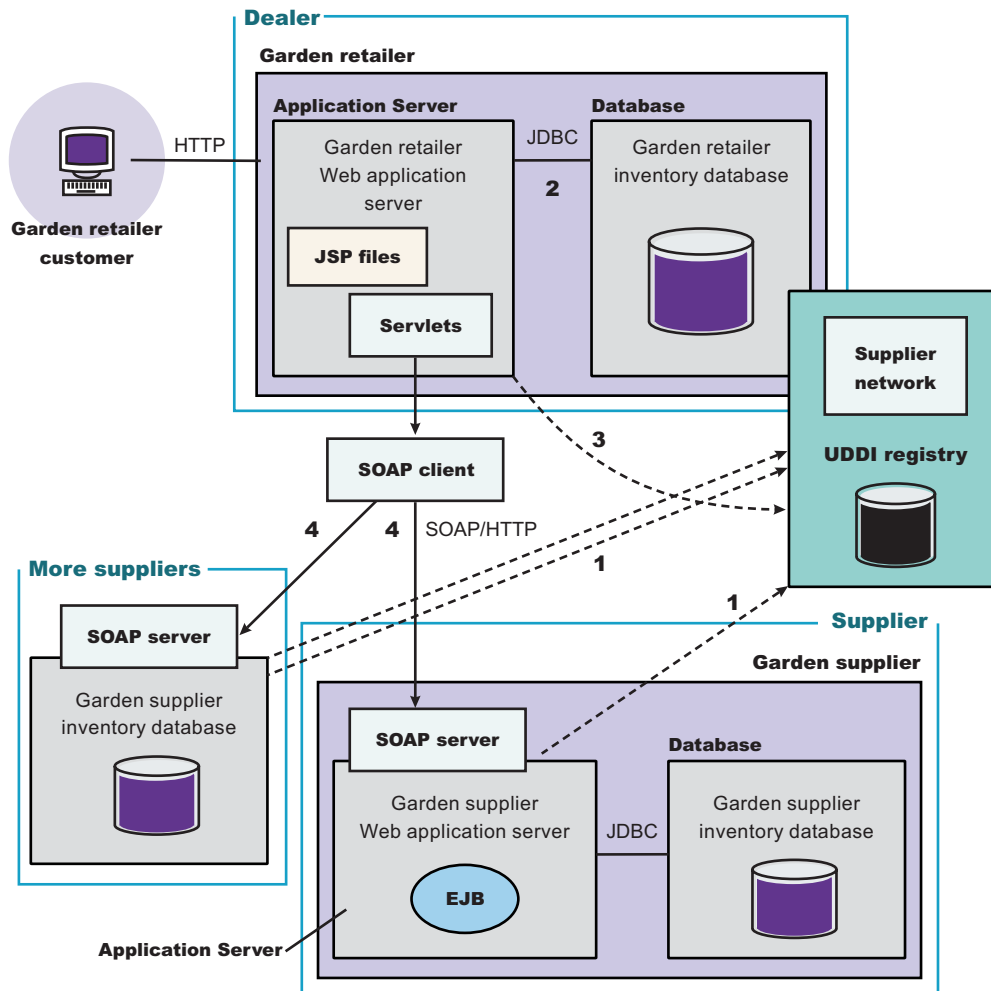
## **Web services scenario: Dynamic inquiry on supplier**

This document describes a scenario in which an online garden supply retailer uses Web services to perform an inventory search of several wholesale suppliers.

In the "static inquiry on suppliers" scenario, the Plants by WebSphere IT staff establishes connections with each supplier separately, and makes changes as suppliers come and go. It would be convenient to query multiple suppliers at the same time, without necessarily even providing a list of particular suppliers to query. Furthermore, Plants by WebSphere managers would like to be able to shop around quickly according to criteria such as the lowest wholesale price or fastest availability.

In a "dynamic inquiry on suppliers" scenario, several plant and garden suppliers have published Web services to a Universal Description, Discovery, and Integration. (UDDI) registry. Suppliers create inventory Web services that use a standard interface. They publish their Web services to the centralized registry. Perhaps the registry has been established by the Plant Wholesalers Association. Or maybe a small Internet company established the buyer-seller site after finding that suppliers and retailers each would pay

a small monthly rate for the convenience of the service.



### How out of stock items are handled

The following events happen when a customer on the Plants by WebSphere site orders an item that is not available according to the Plants by WebSphere inventory.

1. In advance, the suppliers publish their Web services to a UDDI registry for just such an occasion. In this way, they notify inquiring retailers, such as Plants by WebSphere, that their inventories are available to check.

2. Plants by WebSphere checks its own inventory.

The application powering the Web site checks the Plants by WebSphere inventory database. It discovers that the item is not in stock.

3. Plants by WebSphere uses a UDDI4J client to consult the UDDI registry for suppliers whose inventories it can check.

Plants by WebSphere can invoke a Web service that queries the UDDI registry for suppliers. Then in turn, the Web service at the site of each supplier is invoked. The administrator is presented with a list of suppliers from which the requested item is available, perhaps including price and availability data.

4. Plants by WebSphere uses the Web services to check the supplier inventories.

The application invokes a Web services for J2EE client or JAX-RPC client that communicates with a SOAP server at the supplier site to ascertain whether the supplier has the item in stock. The supplier data is sent to Plants by WebSphere.

5. Plants by WebSphere either obtains the out of stock item, or does not.  
If the supplier indicates that the item is in stock, the application powering Plants by WebSphere determines whether to order the item on behalf of the customer. The exchange of data can include checking a price threshold above which Plants by WebSphere will not order the wholesale item. It could also include decision-making information such as how long the supplier will require to deliver the item, or a date on which the manufacturer plans to discontinue the item.
6. Plants by WebSphere notifies its customer of the outcome, as soon as possible.  
If the supplier can be consulted quickly enough, Plants by WebSphere does not have to bother its customer with concerns about availability. It simply confirms that the item is available, as though the item is in stock at Plants by WebSphere. If the supplier inventory temporarily lacks the item, or Plants by WebSphere opted not to order the item from the supplier, Plants by WebSphere will issue an appropriate response to the customer.

### **Web services technologies used in this scenario**

This scenario uses the following Web services technologies.

#### **XML (Extensible Markup Language)**

XML is used to standardize the exchange of data between Plants by WebSphere and its supplier.

#### **Web Services for J2EE**

Web Services for J2EE, also known as JSR-109, defines how J2EE applications create and access Web services.

#### **Java API for Remote Procedure Calls**

JAX-RPC, also known as JSR-101, defines how Java applications access Web services.

#### **WSDL (Web Services Description Language)**

WSDL is used to turn the existing application into a Web service, by acting as the interface between the underlying application and other Web-enabled applications.

#### **SOAP (Simple Object Access Protocol)**

SOAP is the protocol by which the Web service communicates with the supplier over the Internet.

#### **A Universal Description, Discovery and Integration (UDDI) registry**

By publishing their Web services to UDDI, suppliers make them available for Plants by WebSphere and other retailers to discover and reuse. This saves development time, effort and cost, and helps minimize the need to maintain several different implementations of the same application at Plants by WebSphere and various other retailers who need to contact the suppliers for inventory data.

Particular editions of WebSphere Application Server provide a private UDDI registry that can be used in an intranet environment.

---

## **Web services scenario: Cross supplier inquiry**

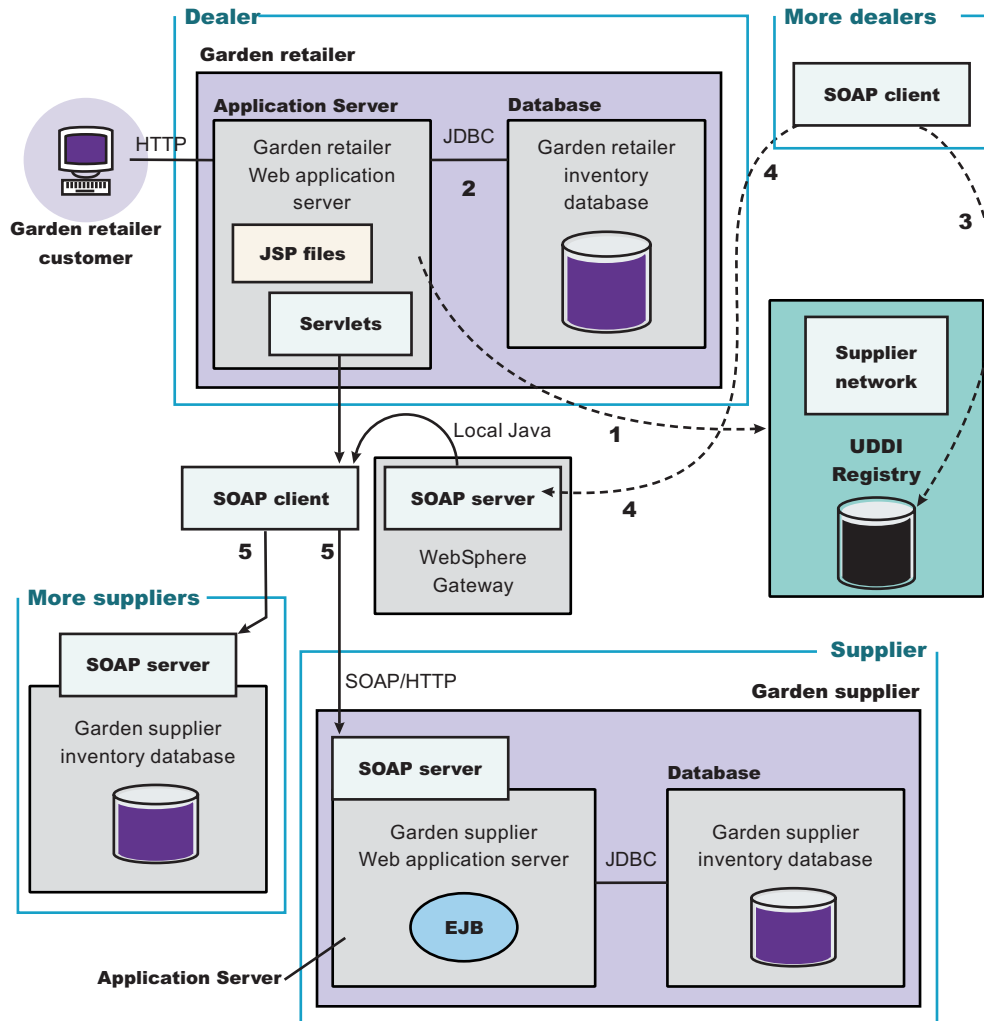
This document describes a scenario in which an online garden supply retailer uses Web services to integrate its inventory system with the inventory systems of other retailers. Also using Web services, the main Internet storefront can check supplier inventories on behalf of itself or other retailers.

The marketers at Plants by WebSphere confirm with market data that people are likely to purchase plants and gardening supplies in tandem with purchases of other goods, such as gardening books. To increase the visibility of Plants by WebSphere, the company arranges with various other merchant sites to include Plants by WebSphere inventory as part of their own.

At one site, Web services and other technologies are used to insert data about Plants by WebSphere items into Web pages that match the look and feel of the rest of the site. When a customer orders a

Plants by WebSphere item at a site other than Plants by WebSphere, the second site relies on the Plants by WebSphere inventory Web service to make sure that the item is in stock, and to query suppliers as needed.

The second site does not have to implement its own Web services to perform the same function as those developed by Plants by WebSphere. The second site might want to implement sophisticated function by creating its own Web service.



### How out of stock items are handled

The following events happen when a customer orders an item from one of the sites that re-sells items from Plants by WebSphere.

1. In advance, Plants by WebSphere publishes its Web service to a public UDDI registry.  
In this way, other retailers are made aware of the inventory Web service available from Plants by WebSphere. In this scenario, Plants by WebSphere probably will enable the Web service to check its own inventory as well as that of suppliers.
2. The re-seller checks the Plants by WebSphere inventory.  
The application powering the Web site checks the Plants by WebSphere inventory database. It discovers that the item is not in stock.
3. The re-seller consults the UDDI registry for suppliers whose inventories it can check.
4. The re-seller uses the Web services to check the supplier inventories.



The application invokes a Web services for J2EE or JAX-RPC SOAP client that communicates with a SOAP server at the supplier site to ascertain whether the supplier has the item in stock. The supplier data is sent to the reseller.

5. The re-seller either obtains the out of stock item, or does not.
6. The re-seller notifies its customer of the outcome, as soon as possible.

### **Web services technologies used in this scenario**

This scenario uses the following Web services technologies.

#### **XML (Extensible Markup Language)**

XML is used to standardize the exchange of data between Plants by WebSphere and its supplier.

#### **WSDL (Web Services Description Language)**

WSDL is used to turn the existing application into a Web service, by acting as the interface between the underlying application and other Web-enabled applications.

#### **SOAP (Simple Object Access Protocol)**

SOAP is the protocol by which the Web service communicates with the supplier over the Internet.

#### **A Universal Description, Discovery and Integration (UDDI) registry**

By publishing their Web services to UDDI, suppliers make them available for Plants by WebSphere and other retailers to discover and reuse. This saves development time, effort and cost, and helps minimize the need to maintain several different implementations of the same application at Plants by WebSphere and various other retailers who need to contact the suppliers for inventory data.

Public UDDI registries are run by a consortium named UDDI Operators Council, which includes IBM, NTT, SAP, and Microsoft.

Particular editions of WebSphere Application Server provide a private UDDI registry that can be used in an intranet environment.

#### **Web Services Invocation Framework (WSIF)**

In addition to publishing SOAP/HTTP bindings to the public UDDI registry for use by other vendors, Plants by WebSphere might also have published to an internal private UDDI registry with additional optimized bindings. A Web service provider such as Plants by WebSphere might offer a SOAP binding for the service and a local Java binding that allows you to treat the local service implementation (a Java class) as a Web service. If the client is deployed in the same environment as the service, then the local Java binding for the service can be used. This provides more efficient communication with the service by making direct Java calls rather than using the SOAP binding.

#### **Web services gateway**

Plants by WebSphere could use a gateway to handle Web service invocations between Internet and Intranet environments. A Web services gateway makes the internal Web service available externally. It takes care of these considerations:

- The transport mechanisms (or channels) on which messages can be carried to and from the service
- The filters (if any) that act upon these incoming and outgoing messages
- The UDDI registries (if any) to which to publish the service
- The levels of security that you want to apply to the service

---

## Chapter 10. Web Services Invocation Framework (WSIF): Enabling Web services

The Web Services Invocation Framework (WSIF) is a WSDL-oriented Java API. You use this API to invoke Web services dynamically, regardless of the service implementation format (for example enterprise bean (EJB)) or the service access mechanism (for example Java Messaging Service (JMS)).

Using WSIF, you can move away from the usual Web services programming model of working directly with the SOAP APIs, towards a model where you interact with representations of the services. You can therefore work with the same programming model regardless of how the service is implemented and accessed.

If you want to know more about the issues that WSIF addresses, see [Goals of WSIF](#).

If you want to know how WSIF addresses these issues, see [An overview of WSIF](#).

To use WSIF, see the following topics:

- Using WSIF to invoke Web services.
- WSIF system management and administration.
- WSIF API.

For more information about working with WSIF, visit the Web sites listed in [Web services: Resources for Learning](#).

---

### Goals of WSIF

SOAP bindings for Web services are part of the WSDL specification, therefore when most developers think of using a Web service, they immediately think of assembling a SOAP message and sending it across the network to the service endpoint, using a SOAP client API. For example: using Apache SOAP the client creates and populates a Call object that encapsulates the service endpoint, the identification of the SOAP operation to invoke, the parameters to send, and so on.

While this process works for SOAP, it is limited in its use as a general model for invoking Web services for the following reasons:

- Web services are more than just SOAP services.
- Tying client code to a particular protocol implementation is restricting.
- Incorporating new bindings into client code is hard.
- Multiple bindings can be used in flexible ways.
- A freer Web services environment enables intermediaries.

The goals of the Web Services Invocation Framework (WSIF) are therefore:

- To give a binding-independent mechanism for Web service invocation.
- To free client code from the complexities of any particular protocol used to access a Web service.
- To enable dynamic selection between multiple bindings to a Web service.
- To help the development of Web service intermediaries.

### WSIF - Web services are more than just SOAP services

You can deploy as a Web service any application that has a WSDL-based description of its functional aspects and access protocols. If you are using the Java 2 platform, Enterprise Edition (J2EE) environment, then the application is available over multiple transports and protocols.

For example, you can take a database-stored procedure, expose it as a stateless session bean, then deploy it into a SOAP router as a SOAP service. At each stage, the fundamental service is the same. All

that changes is the access mechanism: from Java Database Connectivity (JDBC) to Remote Method Invocation over Internet Inter-Orb Protocol (RMI-IIOP) and then to SOAP.

The WSDL specification defines a SOAP binding for Web services, but you can add binding extensions to the WSDL so that, for example, you can offer an enterprise bean as a Web service using RMI-IIOP as the access protocol. You can even treat a single Java class as a Web service, with in-thread Java method invocations as the access protocol. With this broader definition of a Web service, you need a binding-independent mechanism for service invocation.

## **WSIF - Tying client code to a particular protocol implementation is restricting**

If your client code is tightly bound to a client library for a particular protocol implementation, it can become hard to maintain.

For example, if you move from Apache SOAP to Java Messaging Service (JMS) or enterprise bean, the process can take a lot of time and effort. To avoid these problems, you need a protocol implementation-independent mechanism for service invocation.

## **WSIF - Incorporating new bindings into client code is hard**

As is explained in Web services are not just SOAP services, if you want to make an application that uses a custom protocol work as a Web service, you can add extensibility elements to WSDL to define the new bindings. But in practice, achieving this capability is hard. For example you have to design the client APIs to use this protocol. If your application uses just the abstract interface of the Web service, you have to write tools to generate the stubs that enable an abstraction layer. These tasks can take a lot of time and effort. What you need is a service invocation mechanism that allows you to update existing bindings, and to add new bindings.

## **WSIF - Multiple bindings can be used in flexible ways**

Imagine that you have successfully deployed an application that uses a Web service which offers multiple bindings. For example, imagine that you have a SOAP binding for the service and a local Java binding that lets you treat the local service implementation (a Java class) as a Web service.

The local Java binding for the service can only be used if the client is deployed in the same environment as the service. In this case, it is more efficient to communicate with the service by making direct Java calls than by using the SOAP binding.

If your clients could switch the actual binding used based on run-time information, they could choose the most efficient available binding for each situation. To take advantage of Web services that offer multiple bindings, you need a service invocation mechanism that can switch between the available service bindings at run-time, without having to generate or recompile a stub.

## **WSIF - Enabling a freer Web services environment promotes intermediaries**

Web services offer application integrators a loosely-coupled paradigm. In such environments, intermediaries can be very powerful.

Intermediaries are applications that intercept the messages that flow between a service requester and a target Web service, and perform some mediating task (for example logging, high-availability or transformation) before passing on the message. They can be as small as a simple Web service, or as large as the Web services gateway. The Web Services Invocation Framework (WSIF) is designed to make building intermediaries both possible and simple. Using WSIF, intermediaries can add value to the service invocation without needing transport-specific programming.

---

## An overview of WSIF

The Web Services Invocation Framework (WSIF) provides a Java API for invoking Web services, independent of the format of the service or the transport protocol through which it is invoked. This framework addresses all of the issues identified in the goals of WSIF.

WSIF provides the following features:

- An API that provides binding-independent access to any Web service.
- A close relationship with WSDL, so it can invoke any service that you can describe in WSDL.
- A stubless and completely dynamic invocation of a Web service.
- The capability to plug a new or updated implementation of a binding into WSIF at run-time.
- The option to defer the choice of a binding until run-time.

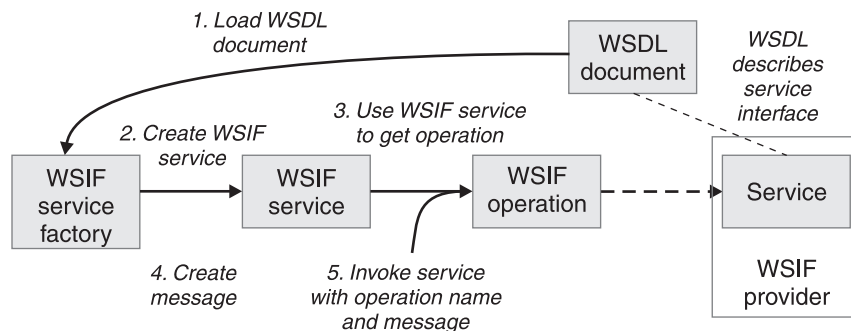
WSIF is designed to work both in an unmanaged environment (stand-alone) and inside a managed container. You can use the Java Naming and Directory Interface (JNDI) to find the WSIF service, or you can use the location described in the WSDL.

For more conceptual information about WSIF and WSDL, see the following topics:

- WSIF and WSDL
- WSIF architecture
- Using WSIF with Web services that offer multiple bindings
- WSIF usage scenarios
- Dynamic invocation

## WSIF architecture

The Web Services Invocation Framework (WSIF) architecture is shown in the figure.



The components of this architecture include:

### WSDL document

The Web service WSDL document contains the location of the Web service. The binding document defines the protocol and format for operations and messages defined by a particular portType.

### WSIF service

The WSIFService interface is responsible for generating an instance of the WSIFOperation interface to use for a particular invocation of a service operation. For more information, see Finding a port factory or service

### WSIF operation

The run-time representation of an operation, called WSIFOperation is responsible for invoking a service based on a particular binding. For more information, see WSIF API reference: Using ports.

### WSIF provider

A WSIF provider is an implementation of a WSDL binding that can run a WSDL operation through a binding-specific protocol. WSIF includes SOAP providers, JMS providers, Java providers and EJB providers. For more information, see Using the WSIF providers.

## Using WSIF with Web services that offer multiple bindings

Using WSIF, a client application can choose dynamically the optimal binding to use for invoking Web service operations.

For example, a Web service might offer a SOAP binding, and also a local Java binding so that you can treat the local service implementation (a Java class) as a Web service. If a client application is deployed in the same environment as the service, then this client can use the local Java binding for the service. This provides more efficient communication between the client and the service by making direct Java calls rather than indirect calls using the SOAP binding.

For more information on how to configure a client to dynamically select between multiple bindings, see [Developing a WSIF service](#).

## WSIF and WSDL

WSDL is the acronym for Web Services Description Language.

In WSDL a service is defined in three distinct sections:

- The **portType**. This section defines the abstract interface offered by the service. A portType defines a set of *operations*. Each operation can be In-Out (request-response), In-Only, Out-Only and Out-In (Solicit-Response). Each operation defines the input and/or output *messages*. A message is defined as a set of *parts*, and each part has a schema-defined type.
- The **binding**. This section defines how to map between the abstract portType and a real service format and protocol. For example the SOAP binding defines the encoding style, the SOAPAction header, the namespace of the body (the targetURI), and so on.
- The **port**. This section defines the actual location (endpoint) of the available service. For example, the HTTP Web address at which a SOAP service is available.

Currently in WSDL, each port has one and only one binding, and each binding has a single portType. But (more importantly) each service (portType) can have multiple ports, each of which represents an alternative location and binding for accessing that service.

The Web Services Invocation Framework (WSIF) follows the semantics of WSDL as much as possible:

- The WSIF dynamic invocation API directly exposes run-time equivalents of the model from WSDL. For example, invocation of an operation involves executing an operation with an input message.
- WSDL has extension points that support the addition of new ports and bindings. This enables WSDL to describe new systems. The equivalent concept in WSIF is a provider, that enables WSIF to understand a class of extensions and thereby to support a new service implementation type.

As a metadata-based invocation framework, WSIF follows the design of the metadata. As WSDL is extended, WSIF is updated to follow.

The implicit and primary type system of WSIF is XML schema. WSIF supports invocation using dynamic proxies, which in turn support Java type systems, but when you use the `WSIFMessage` interface it is your responsibility to populate `WSIFMessage` objects with data based on the XML schema types as defined in the WSDL document. You should define your object types by a canonical and fixed mapping from schema types into the run-time.

For more information on WSDL, see [Web services: Resources for learning](#).

## WSIF usage scenarios

This topic describes two brief scenarios that illustrate the role WSIF plays in the emerging Web services environment.

## Scenario: Redevelopment and redeployment

When you first implement a Web service, you create a simple prototype. When you want to move a prototype Web service into production, you often need to redevelop and redeploy it.

The Web Services Invocation Framework (WSIF) uses the same API calls irrespective of the underlying technologies, therefore if you use WSIF:

- You can reimplement and redeploy your services without changing the client code.
- You can use existing reliable and high-performance infrastructures like Remote Method Invocation over Internet Inter-Orb Protocol (RMI-IIOP) and Java Messaging Service (JMS) without sacrificing the location-independence that the Web service model offers.

## Scenario: Service Flow composition

A service flow typically invokes a Web service, then passes the response from one Web service to the next Web service, perhaps performing some transformation in the middle.

There are two key aspects to this flow that WSIF provides:

- A representation of the service invocation based on the metadata in WSDL.
- The ability to build invocations based solely on the portType, which can therefore be used in any implementation.

For example, imagine that you build a meta-service that uses a number of services to build a process. Initially, several of those services are simple Java bean prototypes that are written and exposed through SOAP, but you plan to reimplement some of them as EJB components, and to out-source others.

If you use SOAP, it ties up multiple threads for every onward invocation, because they pass through the Web server and servlet engine and on to the SOAP router. If you use WSIF to call the beans directly, you get much better performance compared to SOAP and you do not lose access or location transparency. Using WSIF, you can replace the Java bean implementations with EJB implementations without changing the client code. To move some of the Web services from local implementations to external SOAP services, you just update the WSDL.

## Dynamic invocation

For the Web Services Invocation Framework (WSIF), dynamic invocation means providing the following levels of support when invoking Web services:

1. Support for WSDL extensions and bindings that were not known at build time.
2. Support for Web services that were not known at build time.

WSIF supports (1) through the use of providers.

The providers support (2) by using the WSDL description to access the target service.

---

## Using WSIF to invoke Web services

You invoke a Web service dynamically by using the WSIF API directly.

You only specify the location of the WSDL file for the service, the name of the operation to invoke, and any operation arguments. All the information needed to access the Web service (the abstract interface, the binding, and the service endpoint) is available through the WSDL.

This kind of invocation does not require stub classes and does not need a separate compilation cycle.

More information on using the Web Services Invocation Framework (WSIF) to invoke Web services is provided in the following topics:



- Using the WSIF providers.
- Developing a WSIF service.
- Using complex types.
- Using the Java Naming and Directory Interface (JNDI).
- Passing SOAP messages with attachments using WSIF.
- Interacting with the J2EE container in WebSphere Application Server.
- Running WSIF as a client.

## Using the WSIF providers

A Web Services Invocation Framework (WSIF) provider is an implementation of a WSDL binding that can run a WSDL operation through a binding-specific protocol.

Providers implement the interface between the WSIF API and the actual implementation of a service. Providers are pluggable within the WSIF framework, and are registered according to the namespace of the WSDL extension that they implement. Some providers use the Java 2 platform, Enterprise Edition (J2EE) programming model to utilize J2EE services. If a provider is available, but its required class libraries are not, then the provider is disabled.

WebSphere Application Server includes the following WSIF providers:

- SOAP (over HTTP) provider.
- JMS providers (SOAP over JMS, and native JMS).
- Java provider.
- EJB provider.

## Using the SOAP provider

The SOAP provider allows WSIF stubs and dynamic clients to invoke SOAP services.

The Web Services Invocation Framework (WSIF) SOAP provider supports SOAP 1.1 over HTTP. The SOAP provider uses Apache SOAP 2.3 for parsing and creating SOAP messages, but it is not limited to invoking services from Apache SOAP.

The WSIF SOAP provider supports:

- SOAP-ENC encoding.
- RPC style and Document style SOAP messages.
- SOAP messages with attachments.

The SOAP provider is not transactional.

If you have a Web service that you expect multiple clients to use connecting over SOAP, then before you deploy the service you must set up your application deployment descriptor file `dds.xml` to handle multiple connections correctly. For more information, see WSIF troubleshooting tips.

For an example of the sort of code changes that need to be made in the WSDL file for a SOAP provider, see the following topics:

- The SOAP over JMS provider - writing the WSDL extension.
- SOAP messages with attachments - Writing the WSDL extensions.

## Using the JMS providers

The JMS providers enable a WSIF service to be invoked through JMS.

The Java Messaging Service (JMS) is an API for transport technology. The mapping to a JMS destination is defined during deployment and maintained by the container.

The JMS destination endpoint for a Web service can be realized in any of the following ways:

- The JMS destination for the queue can be the Web service implementation.



- The JMS destination can be (but is not required to be) associated with a message-driven bean by the EJB container, thereby allowing the message-driven bean to be the Web service implementation.
- (For SOAP over JMS) The JMS destination can unwrap the JMS message and route the SOAP message to a Web service that is implemented as a stateless session bean.

The JMS destination endpoint must respect the interaction model expected by the client and defined by the WSDL. It must return a response if one is required.

When the JMS destination endpoint creates the JMS response message the following rules must be followed:

- The response message must be sent to `JMSReplyTo` from the incoming request.
- The `JMSCorrelationID` value of the response message must be set to the `JMSMessageID` value from the request message.
- The response must be sent with a `deliveryMode` value equal to the `JMSDeliveryMode` value of the request message.
- The response must be sent with a `priority` value equal to the `JMSPriority` value of the request message.
- The `timeToLive/JMSExpiration` value must be set to a value that equals the `JMSExpiration` value of the request message.

The client does not see any of these headers. The container receives the JMS message and (for SOAP over JMS) removes the SOAP message to send to the client.

See also the following topics:

- Using the SOAP over JMS provider
- Using the native JMS provider
- The JMS providers - Configuring the client and server

#### ***Using the SOAP over JMS provider:***

For information on working with the Java Messaging Service (JMS) API, see [Using the JMS providers](#).

The SOAP message, including the SOAP envelope, is wrapped with a JMS message and put on the appropriate queue. The container receives the JMS message and removes the SOAP message to send to the client.

For detailed implementation information, see the following topics:

- The SOAP over JMS provider - writing the WSDL extension
- The JMS providers - Configuring the client and server

#### *The SOAP over JMS provider - Writing the WSDL extension:*

If a SOAP message contains only XML, then it can be carried on the Java Messaging Service (JMS) transport with the JMS message body type **TextMessage**.

The WSDL binding extension for SOAP over JMS varies only slightly from the SOAP over HTTP binding.

#### **Selecting the SOAP over JMS binding**

You set the `transport` attribute of the `<soap:binding>` tag to indicate that JMS is used. If you also set the `style` attribute to `rpc` (Remote Procedure Call), then the Web Services Invocation Framework (WSIF) assumes that an operation is invoked on the Web service endpoint:

```
<soap:binding style="rpc" transport="http://schemas.xmlsoap.org/soap/jms"/>
```

#### **Setting the JMS address**

For SOAP over JMS, the `<wsdl:port>` tag must contain a `<jms:address>` element. This element provides the information required for a client to connect correctly to the Web service using the

JMS programming model. Typically, it is the stubs generated to support the SOAP over JMS binding that act as the JMS client. Alternatively, the Web service client can use the JMS programming model directly.

The `<jms:address>` element takes this form:

```
<jms:address
 destinationStyle="queue"
 jmsVendorURI="http://ibm.com/ns/mqseries"?
 initialContextFactory="com.ibm.NamingFactory"?
 jndiProviderURL="iiop://something:900/wherever"?
 jndiConnectionFactoryName="orange"
 jndiDestinationName="fred"
/>
```

where attributes marked with a question mark (?) are optional.

The optional `jmsVendorURI` attribute is a string that uniquely identifies the JMS implementation. WSIF ignores this URI, which is used by the client developer and perhaps the client implementation to determine if it has access to the correct JMS provider in the client run-time.

The optional attributes `initialContextFactory` and `jndiProviderURL` can only be omitted if the run-time has a default Java Naming and Directory Interface (JNDI) provider configured.

The `jndiConnectionFactoryName` attribute gives the name of a JMS `ConnectionFactory` object, which can be looked up within the JNDI context given by the `jndiContext` attribute. This `ConnectionFactory` object is used to create a JMS connection to the JMS provider instance that owns the queue. In a simple configuration, the same `ConnectionFactory` object is used by the server message listener and by the clients. However the server and the clients can use different `ConnectionFactory` objects, provided that they all create connections to the same JMS provider instance.

### Setting the JMS headers and properties

You use the `<jms:property>` tag to set the JMS headers and properties. This tag maps either a message part, or a literal value, into a JMS property:

```
<jms:property name="Priority" {part="requestPriority" | value="fixedValue"}/>
```

If the `<jms:property>` has a literal value, then it can also be nested within the `<jms:address>` tag:

```
<jms:property name="Priority" value="fixedValue" />
```

This form of the `<jms:property>` tag is also used in the native JMS binding.

Here is an example of a WSDL that defines a SOAP over JMS binding:

```
<!-- Example: SOAP over JMS Text Message -->
<?xml version="1.0" encoding="UTF-8"?>
<wsdl:definitions
 name="StockQuoteInterfaceDefinitions"
 targetNamespace="urn:StockQuoteInterface"
 xmlns:tns="urn:StockQuoteInterface"
 xmlns:xsd="http://www.w3.org/2000/10/XMLSchema"
 xmlns:soap="http://schemas.xmlsoap.org/wsdl/soap/"
 xmlns:wsdl="http://schemas.xmlsoap.org/wsdl/">
 <wsdl:message name="GetQuoteInput">
 <part name="symbol" type="xsd:string"/>
 </wsdl:message>
 <wsdl:message name="GetQuoteOutput">
 <part name="value" type="xsd:float"/>
 </wsdl:message>
```

```

<wsdl:portType name="StockQuoteInterface">
 <wsdl:operation name="GetQuote">
 <wsdl:input message="tns:GetQuoteInput"/>
 <wsdl:output message="tns:GetQuoteOutput"/>
 </wsdl:operation>
</wsdl:portType>

<wsdl:binding name="StockQuoteSoapJMSBinding" type="tns:StockQuoteInterface">
 <soap:binding style="rpc"
 transport="http://schemas.xmlsoap.org/soap/jms"/>
 <wsdl:operation name="GetQuote">
 <soap:operation soapAction="urn:StockQuoteInterface#GetQuote"/>
 <wsdl:input>
 <soap:body use="encoded" namespace="urn:StockQuoteService"
 encodingStyle="http://schemas.xmlsoap.org/soap/encoding"/>
 </wsdl:input>
 <wsdl:output>
 <soap:body use="encoded" namespace="urn:StockQuoteService"
 encodingStyle="http://schemas.xmlsoap.org/soap/encoding"/>
 </wsdl:output>
 </wsdl:operation>
</wsdl:binding>
<wsdl:service name="StockQuoteService">
 <wsdl:port name="StockQuoteServicePort"
 binding="sqi:StockQuoteSoapJMSBinding">
 <jms:address destinationStyle="queue"
 jndiConnectionFactoryName="myQCF"
 jndiDestinationName="myQ"
 initialContextFactory="com.ibm.NamingFactory"
 jndiProviderURL="iiop://something:900/" />
 </wsdl:port>
</wsdl:service>
</wsdl:definitions>

```

### **Using the native JMS provider:**

Using the native JMS provider, WSIF clients can treat a JMS destination as a Web service.

For information on working with the Java Messaging Service (JMS) API, see [Using the JMS providers](#).

For detailed implementation information, see the following topics:

- [The native JMS provider - Writing the WSDL extension](#)
- [The JMS providers - Configuring the client and server](#)

#### *The native JMS provider - Writing the WSDL extension:*

The WSDL extensions for the Java Messaging Service (JMS) are identified with the namespace prefix `jms`. For example, `<jms:binding>`.

### **Operations**

The supported operations are either one-way operations (send for JMS point-to-point messaging, or publish for JMS publish and subscribe messaging) or request-response operations (send and receive for JMS point-to-point messaging). The WSDL operations therefore specify either an input message only, or an input and an output message.

### **Fault messages**

Operations that describe message interfaces with a native JMS binding do not have fault messages. No assumptions are made about the message schema or the semantics of message properties, therefore no distinction can be made between output and fault messages.

### **Setting the JMS message body type**

You use the `<jms:binding>` extension to specify the JMS message body type:

```
<wsdl:binding ... >
 <jms:binding type="messageBodyType" />
 ...
</wsdl:binding>
```

where *messageBodyType* is either `ObjectMessage` or `TextMessage`.

### Specifying the parts to use for the input and output messages

For JMS text messages and JMS object messages created from one or more WSDL message parts, you use the `<jms:input>` and `<jms:output>` extensions to specify the message parts to use for the JMS messages:

```
<wsdl:input ... >
 <jms:input parts="part1 part2 ..." />
</wsdl:input>

<wsdl:output ... >
 <jms:output parts="part1 part2 ..." />
</wsdl:output>
```

In the next example, the WSDL message has just one part that contains the complete message body. This message body might result from a mapping of some other representation (see **Mapping data types**).

```
<wsdl:input ... >
 <jms:input parts="part1" />
</wsdl:input>
```

If no parts are defined, then all the message parts are used.

### Mapping data types

You use the `<format>` extensions to map data types:

```
<wsdl:binding ... >
 <jms:binding type="..." />

 <format:typeMapping encoding="Java" style="Java">
 <format:typeMap typeName="..." formatType="targetType"/>
 </format:typemapping>
 ...
</wsdl:binding>
```

The value of *targetType* is dependent on the JMS message body type (see **Setting the JMS message body type**). For JMS object messages, the target data type implements the `java.io.Serializable` class. For JMS text messages, the target data type is always `java.lang.String`.

The `<format>` extensions are also used in other bindings that deal with Java interfaces.

### Setting the JMS headers and properties

JMS does not make assumptions about message headers. For example, if the JMS provider is `MQSeries` then each JMS message carries an `RFH2` header. However you can access data in this message header indirectly, by getting and setting JMS message properties.

When you want your application to pass a property into the Web Services Invocation Framework (WSIF) as a part on the WSIF message, you use a `<jms:property>` tag. When you want to hard code an actual property value into the WSDL, you use a `<jms:propertyValue>` tag. The `<jms:propertyValue>` tag contains a specification of a literal value and its associated XML schema type.

You can specify `<jms:property>` and `<jms:propertyValue>` extensions within the `<wsdl:input>` tag in the binding operation, and also within the `<jms:address>` tag. For the `<wsdl:output>` tag in the

binding operation, you can only specify the `<jms:property>` extension. Property values that are set in the `<jms:property>` tag take precedence over values set in the `<<jms:propertyValue>` tag, and property values that are set in the binding operation (in the `<input>` and `<output>` tags) take precedence over values set in the `<jms:address>` tag.

Here is an example of the `<jms:property>` and `<jms:propertyValue>` tags nested within the `<input>` and `<output>` tags:

```
<wsdl:input ... >
 <jms:property name="propertyName" part="partName" />
 <jms:propertyValue name="propertyName"
 type="xsdType" value="actualValue" />
</wsdl:input>
<wsdl:output ... >
 <jms:property name="propertyName" part="partName" />
</wsdl:output>
```

where *propertyName* identifies the JMS property that is associated with the header field, and *partName* identifies the message part that is associated with the property.

The JMS property identified by *propertyName* can be user-defined, or it can be one of the following predefined JMS message header fields:

Value	Java type
JMSMessageId	java.lang.String
JMSTimeStamp	long
JMSCorrelationId	byte [ ] or java.lang.String
JMSReplyTo	javax.jms.Destination
JMSDestination	javax.jms.Destination
JMSDeliveryMode	int
JMSRedelivered	boolean
JMSType	java.lang.String
JMSExpiration	long

See the JMS specification for restrictions that apply when setting JMS header field values. Attempts to set restricted values are ignored.

For application-defined JMS message properties, the Java types used in the native JMS binding implementation (used for calls to the corresponding JMS methods) are derived from the XML schema type in the abstract interface (`<wsdl:part>` tag), and from the type mapping information in the format binding (`<format:typemap>` tag).

## Handling transactions

Independent of other JMS properties, the asynchronous processing of request-response operations has implications for callers running in a transaction scope. The send request part and the receive response part are separated into two transactions, because the send needs to be committed in order for the request message to become visible. Implementations that process WSDL for asynchronous request-response operations (such as WSIF) must therefore take the following additional actions:

- They must ensure that the send request transaction returns a correlation ID to the user, and provides a **callback** with which users can pass in the response message to process the receive response transaction.
- They might implement their own response message “listener” in order to recognize the arrival of response messages, and to manage the correlation to the request message.

The JMS text message contains a **java.lang.String**. In this example, the WSDL message contains only one part that represents the whole message body:

```

<!-- Example 1: JMS Text Message -->

<wsdl:definitions ... >

 <!-- simple or complex types for input and output message -->
 <wsdl:types> ... </wsdl:types>

 <wsdl:message name="JmsOperationRequest"> ... </wsdl:message>
 <wsdl:message name="JmsOperationResponse"> ... </wsdl:message>

 <wsdl:portType name="JmsPortType">
 <wsdl:operation name="JmsOperation">
 <wsdl:input name="Request"
 message="tns:JmsOperationRequest"/>
 <wsdl:output name="Response"
 message="tns:JmsOperationResponse"/>
 </wsdl:operation>
 </wsdl:portType>

 <wsdl:binding name="JmsBinding" type="JmsPortType">
 <jms:binding type="TextMessage" />

 <format:typemapping style="Java" encoding="Java">
 <format:typemap name="xsd:String" formatType="String" />
 </format:typemapping>

 <wsdl:operation name="JmsOperation">
 <wsdl:input message="JmsOperationRequest">
 <jms:input parts="requestMessageBody" />
 </wsdl:input>
 <wsdl:output message="JmsOperationResponse">
 <jms:output parts="responseMessageBody" />
 </wsdl:output>
 </wsdl:operation>
 </wsdl:binding>

 <wsdl:service name="JmsService">
 <wsdl:port name="JmsPort" binding="JmsBinding">
 <jms:address destinationStyle="queue"
 jndiConnectionFactoryName="myQCF"
 jndiDestinationName="myDestination"/>
 </wsdl:port>
 </wsdl:service>

</wsdl:definitions>

```

As an extension to the previous JMS message example, the following example WSDL describes a request-response operation in which specific JMS property values of the request and response message are set for the request message and retrieved from the response message.

The JMS properties in the request message are set according to the values in the input message. Likewise, selected JMS properties of the response message are copied to the corresponding values of the output message. The direction of the mapping is determined by the appearance of the `<jms:property>` tag in the input or output section, respectively.

```

<!-- Example 2: JMS Message with JMS Properties -->
<wsdl:definitions ... >

 <!-- simple or complex types for input and output message -->
 <wsdl:types> ... </wsdl:types>

 <wsdl:message name="JmsOperationRequest">
 <wsdl:part name="myInt" type="xsd:int"/>
 ...
 </wsdl:message>

 <wsdl:message name="JmsOperationResponse">
 <wsdl:part name="myString" type="xsd:String"/>
 ...
 </wsdl:message>

 <wsdl:portType name="JmsPortType">
 <wsdl:operation name="JmsOperation">
 <wsdl:input name="Request"
 message="tns:JmsOperationRequest"/>
 <wsdl:output name="Response"
 message="tns:JmsOperationResponse"/>
 </wsdl:operation>
 </wsdl:portType>

 <wsdl:binding name="JmsBinding" type="JmsPortType">
 <!-- the JMS message type may be any of the above -->
 <jms:binding type="..." />

 <format:typemapping style="Java" encoding="Java">
 <format:typemap name="xsd:int" formatType="int" />
 ...
 </format:typemapping>

 <wsdl:operation name="JmsOperation">
 <wsdl:input message="JmsOperationRequest">
 <jms:property message="tns:JmsOperationRequest" parts="myInt" />
 <jms:propertyValue name="myLiteralString"
 type="xsd:string" value="Hello World" />
 ...
 </wsdl:input>
 <wsdl:output message="JmsOperationResponse">
 <jms:property message="tns:JmsOperationResponse" parts="myString" />
 ...
 </wsdl:output>
 </wsdl:operation>
 </wsdl:binding>

 <wsdl:service name="JmsService">
 <wsdl:port name="JmsPort" binding="JmsBinding">
 <jms:address destinationStyle="queue"
 jndiConnectionFactoryName="myQCF"
 jndiDestinationName="myDestination"/>
 </wsdl:port>
 </wsdl:service>
</wsdl:definitions>

```

### ***The JMS providers - Configuring the client and server:***

This topic assumes that you installed a Java Messaging Service (JMS) provider when you installed WebSphere Application Server (either the JMS provider that is embedded in WebSphere Application Server, or another provider such as WebSphere MQ). If not, install one now as described in Installing and configuring a JMS provider.



To enable a service to be invoked through JMS by a Web Services Invocation Framework (WSIF) client application, complete the following steps:

1. Use the administrative console to create and configure a queue connection factory and a queue destination as described in *Configuring JMS provider resources*.
2. Use the administrative console to add the new queue destination to the list of JMS Server destination names for your application server as described in *Managing JMS servers on Application Server* (not part of a deployment manager cell) or *Managing JMS servers in a deployment manager cell*. Ensure that the Initial State is started.
3. Put the JNDI names of the queue destination and queue connection factory, as well as your JNDI configuration, in the WSDL file.

You should also understand the specific ways in which WSIF interacts with JMS:

- Only input JMS properties are supported.
- WSIF needs two queues when invoking an operation: one for the request message and one for the reply. The replyTo queue is by default a temporary queue, which WSIF creates on behalf of the application. You can specify a permanent queue by setting the JMSReplyTo property to the JNDI name of a queue.
- WSIF uses the default values for properties set by the JMS implementation. However in MQSeries and in some other JMS implementations, messages are persistent by default, and the default temporary queue is of type temporary dynamic and cannot have persistent messages written to it. Therefore your JMS listener can fail to write a persistent response message to the temporary replyTo queue.

**Note:** If you are using MQSeries, you need to create a temporary model queue that is of type permanent dynamic, then pass this model as the tempmodel of your queue connection factory. This will ensure that persistent messages are written to a temporary replyTo queue that is of type permanent dynamic.

## Using the Java provider

Using the WSIF Java provider, WSIF can invoke Java code.

This means that, in a thin-client environment such as a Java Virtual Machine (JVM) or Tomcat test run-time, you can define shortcuts to local Java programs.

The Web Services Invocation Framework (WSIF) Java provider is not intended for use in a Java 2 platform, Enterprise Edition (J2EE) environment. There is a difference between a client using the WSIF Java provider to invoke a Java component, and implementing a Web service as a Java component on the server side.

The Java binding exploits the format binding for type mapping. Using the format binding, your WSDL can define the mapping between XML schema types and Java types.

The Java provider requires that the targeted Java classes reside in the class path of the client. The Java method is invoked synchronously, in-process, in-thread, with the current thread and Object Request Broker (ORB) contexts.

The Java provider is not transactional.

For examples of the code changes that need to be made in the WSDL file, see *The Java provider - Writing the WSDL extension*.

### ***The Java provider - Writing the WSDL extension:***

The Java provider supports the invocation of a method on a local Java object.

To use the Java provider, you need the following binding specified in the WSDL:

```

<!-- Java binding -->
<binding >
 <java:binding />
 <format:typeMapping style="Java" encoding="Java"/>?
 <format:typeMap name="qname" formatType="nmtoken"/>*
 </format:typeMapping>
 <operation>*
 <java:operation
 methodName="nmtoken"
 parameterOrder="nmtoken"
 returnPart="nmtoken"?
 methodType="instance|constructor" />
 <input name="nmtoken"? />?
 <output name="nmtoken"? />?
 <fault name="nmtoken"? />?
 </operation>
</binding>

```

In this example:

- A question mark (?) means optional, and an asterisk (\*) means 0 or more.
- The name attribute of the <format:typeMap> element is a qualified name of a simple or complex type used by one of the Java operations.
- The formatType attribute of the <format:typeMap> element is the fully qualified class name for the Java class to which the element specified by name maps.
- The methodName attribute of the <java:operation> element is the name of the method on the Java object that is called by the operation.
- The parameterOrder attribute of the <java:operation> element contains a white space-separated list of part names that define the order in which they are passed to the Java object method.
- The methodType attribute of the <java:operation> element must be set to either instance or constructor. The value specifies whether the method that is invoked on the object is an instance method or a constructor for the object.

In the next example, the className attribute of the <java:address> element specifies the fully qualified class name of the object containing the method to invoke:

```

<service ... >
 <port>*
 <java:address
 className="nmtoken"/>
 </port>
</service>

```

## Using the EJB provider

Using the EJB provider, WSIF clients can invoke enterprise beans.

The EJB client JAR file must be available in the client run-time with the current provider. The enterprise bean is invoked using normal EJB invocation methods, using Remote Method Invocation over Internet Inter-Orb Protocol (RMI-IIOP), with the current security and transaction contexts. If the EJB provider is invoked within a transaction, the transaction is passed to the onward service and the standard EJB transaction attribute applies.

If there are multiple implementations of the service, it is up to the service providers to make sure that every implementation offers the same semantics. For example, in the case of transactions, the bean deployer must specify TX\_REQUIRES\_NEW to force a new transaction.

For examples of the code changes that need to be made in the WSDL file, see The EJB provider - Writing the WSDL.

### *The EJB provider - Writing the WSDL extension:*

The EJB provider supports the invocation of an enterprise bean through Remote Method Invocation over Internet Inter-Orb Protocol (RMI-IIOP).

To use the EJB provider, you need the following binding specified in the WSDL:

```
<!-- EJB binding -->
<binding >
 <ejb:binding />
 <format:typeMapping style="Java" encoding="Java"/>?
 <format:typeMap name="qname" formatType="nmtoken"/>*
 </format:typeMapping>
 <operation>*
 <ejb:operation
 methodName="nmtoken"
 parameterOrder="nmtoken"
 returnPart="nmtoken"?
 interface="remote|home" />
 <input name="nmtoken"? />?
 <output name="nmtoken"? />?
 <fault name="nmtoken"? />?
 </operation>
</binding>
```

In this example:

- A question mark (?) means optional, and an asterisk (\*) means 0 or more.
- The name attribute of the <format:typeMap> element is a qualified name of a simple or complex type used by one of the EJB operations.
- The formatType attribute of the <format:typeMap> element is the fully qualified class name for the Java class to which the element specified by name maps.
- The methodName attribute of the <ejb:operation> element is the name of the method on the enterprise bean that is called by the operation.
- The parameterOrder attribute of the <ejb:operation> element contains a white space-separated list of part names that define the order in which they are passed to the EJB method.
- The interface attribute of the <ejb:operation> element must be set to either remote or home. The value specifies the interface of the enterprise bean on which the method named by the methodName attribute is accessible.

In the next example:

- The className attribute of the <ejb:address> element specifies the fully qualified class name of the home interface class of the enterprise bean.
- The jndiName attribute of the <ejb:address> element specifies the full Java Naming and Directory Interface (JNDI) name that is used to look up the enterprise bean.
- The initialContextFactory attribute of the <ejb:address> element is optional and specifies the initial context factory class.
- The jndiProviderURL attribute of the <ejb:address> element is optional and specifies the JNDI provider Web address.

```
<service ... >
 <port>*
 <ejb:address
 className="nmtoken"
 jndiName="nmtoken"
 initialContextFactory="nmtoken" ?
 jndiProviderURL="nmtoken" ? />
 </port>
</service>
```

## Developing a WSIF service

A Web Services Invocation Framework (WSIF) service is a Web service that uses WSIF.

To develop a WSIF service, develop the Web service (or use an existing Web service), then develop the WSIF client based on the WSDL document for that Web service.

There are also two pre-built WSIF Samples available for download from the Samples Central page of the DeveloperWorks WebSphere Web site:

- The Address Book Sample.
- The Stock Quote Sample.

For more information on using the pre-built Samples, see the documentation that is included in the download package.

To develop a WSIF service, complete the following steps:

1. Develop the Web service.

Use Web services tools to discover, create, and publish the Web service. You can develop Java bean, enterprise bean, and URL Web services. You can use Web service tools to create skeleton Java code and a sample application from a WSDL document. For example, an enterprise bean can be offered as a Web service, using Remote Method Invocation over Internet Inter-Orb Protocol (RMI-IIOP) as the access protocol. Or you can use a Java class as a Web service, with native Java invocations as the access protocol.

You can use the WebSphere Studio Application Developer to create a Web service from a Java application, as described in its StockQuote service tutorial. The Java application that you use in this scenario returns the last trading price from the Internet Web site [www.xmltoday.com](http://www.xmltoday.com), given a stock symbol. Using the Web service wizard, you generate a binding WSDL document named `StockQuoteService-binding.wsdl` and a service WSDL document named `StockQuoteService-service.wsdl` from the `StockQuoteService.java` bean. You then deploy the Web service to a Web server, generate a client proxy to the Web service, and generate a sample application that accesses the `StockQuoteService` through the client proxy. You test the `StockQuote` Web service, publish it using the IBM UDDI Explorer, and then discover the `StockQuote` Web service in the IBM UDDI Test Registry.

2. Develop the WSIF client. The information you need to develop a WSIF client is provided in the following topics:

- [Developing the WSIF client - the Address Book Sample](#) gives example code to show how you define a Web service in WSDL.
- [Using the WSIF providers](#) describes the available providers, and gives example code of how their WSDL extensions are coded.
- [WSIF API](#) defines the main interfaces that your client uses to support the invocation of Web services defined in WSDL.

The Address Book Sample is written for synchronous interaction. If you are using a JMS provider, your WSIF client might need to act asynchronously. WSIF provides two main features that meet this requirement:

- A **correlation service** that assigns identifiers to messages so that the request can match up with the (eventual) response.
- A **response handler** that picks up the response from the Web service at a later time.

For more information, see the WSIF API topic [WSIFOperation - Asynchronous interactions](#) reference.

## Developing the WSIF client - the Address Book Sample

The code fragments in this topic show you how to use the Web Services Invocation Framework (WSIF) API to invoke the AddressBook Sample Web service dynamically.

This is example code for dynamic invocation of the AddressBook sample Web service using WSIF:

```
try {
 String wsdlLocation="clients/addressbook/AddressBookSample.wsdl";

 // The starting point for any dynamic invocation using wsif is a
 // WSIFServiceFactory. We create ourselves one via the newInstance
 // method.
```

```

WSIFServiceFactory factory = WSIFServiceFactory.newInstance();

// Once we have a factory, we can use it to create a WSIFService object
// corresponding to the AddressBookService service in the wsdl file.
// Note: since we only have one service defined in the wsdl file, we
// do not need to use the namespace and name of the service and can pass
// null instead. This also applies to the port type, although values have
// been used below for illustrative purposes.
WSIFService service = factory.getService(
 wsdlLocation, // location of the wsdl file
 null, // service namespace
 null, // service name
 "http://www.ibm.com/namespace/wsif/samples/ab", // port type namespace
 "AddressBookPT" // port type name
);

// The AddressBook.wsdl file contains the definitions for two complexType
// elements within the schema element. We will now map these complexTypes
// to Java classes. These mappings are used by the Apache SOAP provider
service.mapType(
 new javax.xml.namespace.QName(
 "http://www.ibm.com/namespace/wsif/samples/ab/types",
 "address"),
 Class.forName("com.ibm.www.namespace.wsif.samples.ab.types.WSIFAddress"));
service.mapType(
 new javax.xml.namespace.QName(
 "http://www.ibm.com/namespace/wsif/samples/ab/types",
 "phone"),
 Class.forName("com.ibm.www.namespace.wsif.samples.ab.types.WSIFPhone"));
// We now have a WSIFService object. The next step is to create a WSIFPort
// object for the port we wish to use. The getPort(String portName) method
// allows us to generate a WSIFPort from the port name.
WSIFPort port = null;

if (portName != null) {
 port = service.getPort(portName);
}
if (port == null) {
 // If no port name was specified, attempt to create a WSIFPort from
 // the available ports for the port type specified on the service
 port = getPortFromAvailablePortNames(service);
}

// Once we have a WSIFPort, we can create an operation. We are going to execute
// the addEntry operation and therefore we attempt to create a WSIFOperation
// corresponding to it. The addEntry operation is overloaded in the wsdl ie.
// there are two versions of it, each taking different parameters (parts).
// This overloading requires that we specify the input and output message
// names for the operation in the createOperation method so that the correct
// operation can be resolved.
// Since the addEntry operation has no output message, we use null for its name.
WSIFOperation operation =
 port.createOperation("addEntry", "AddEntryWholeNameRequest", null);

// Create messages to use in the execution of the operation. This should
// be done by invoking the createXXXXXMessage methods on the WSIFOperation.
WSIFMessage inputMessage = operation.createInputMessage();
WSIFMessage outputMessage = operation.createOutputMessage();
WSIFMessage faultMessage = operation.createFaultMessage();

// Create a name and address to add to the addressbook
String nameToAdd="Chris P. Bacon";
WSIFAddress addressToAdd =
 new WSIFAddress (1,
 "The Waterfront",
 "Some City",
 "NY",

```

```

 47907,
 new WSIFPhone (765, "494", "4900"));

// Add the name and address to the input message
inputMessage.setObjectPart("name", nameToAdd);
inputMessage.setObjectPart("address", addressToAdd);

// Execute the operation, obtaining a flag to indicate its success
boolean operationSucceeded =
 operation.executeRequestResponseOperation(
 inputMessage,
 outputMessage,
 faultMessage);

if (operationSucceeded) {
 System.out.println("Successfully added name and address to addressbook\n");
} else {
 System.out.println("Failed to add name and address to addressbook");
}

// Start from fresh
operation = null;
inputMessage = null;
outputMessage = null;
faultMessage = null;

// This time we will lookup an address from the addressbook.
// The getAddressFromName operation is not overloaded in the
// wsdl and therefore we can simply specify the operation name
// without any input or output message names.
operation = port.createOperation("getAddressFromName");

// Create the messages
inputMessage = operation.createInputMessage();
outputMessage = operation.createOutputMessage();
faultMessage = operation.createFaultMessage();

// Set the name to find in the addressbook
String nameToLookup="Chris P. Bacon";
inputMessage.setObjectPart("name", nameToLookup);

// Execute the operation
operationSucceeded =
 operation.executeRequestResponseOperation(
 inputMessage,
 outputMessage,
 faultMessage);

if (operationSucceeded) {
 System.out.println("Successful lookup of name '"+nameToLookup+"' in addressbook");

 // We can obtain the address that was found by querying the output message
 WSIFAddress addressFound = (WSIFAddress) outputMessage.getObjectPart("address");
 System.out.println("The address found was:");
 System.out.println(addressFound);
} else {
 System.out.println("Failed to lookup name in addressbook");
}

} catch (Exception e) {
 System.out.println("An exception occurred when running the sample.");
 e.printStackTrace();
}
}

```

The preceding code refers to the following Sample method:

```

WSIFPort getPortFromAvailablePortNames(WSIFService service)
 throws WSIFException {
 String portChosen = null;

 // Obtain a list of the available port names for the service
 Iterator it = service.getAvailablePortNames();
 {
 System.out.println("Available ports for the service are: ");
 while (it.hasNext()) {
 String nextPort = (String) it.next();
 if (portChosen == null)
 portChosen = nextPort;
 System.out.println(" - " + nextPort);
 }
 }
 if (portChosen == null) {
 throw new WSIFException("No ports found for the service!");
 }
 System.out.println("Using port " + portChosen + "\n");

 // An alternative way of specifying the port to use on the service
 // is to use the setPreferredPort method. Once a preferred port has
 // been set on the service, a WSIFPort can be obtained via getPort
 // (no arguments). If a preferred port has not been set and more than
 // one port is available for the port type specified in the WSIFService,
 // an exception is thrown.
 service.setPreferredPort(portChosen);
 WSIFPort port = service.getPort();
 return port;
}

```

The Web service uses the following classes:

#### **WSIFAddress:**

```

public class WSIFAddress implements Serializable {

 //instance variables
 private int streetNum;
 private java.lang.String streetName;
 private java.lang.String city;
 private java.lang.String state;
 private int zip;
 private WSIFPhone phoneNumber;

 //constructors
 public WSIFAddress () { }

 public WSIFAddress (int streetNum,
 java.lang.String streetName,
 java.lang.String city,
 java.lang.String state,
 int zip,
 WSIFPhone phoneNumber) {
 this.streetNum = streetNum;
 this.streetName = streetName;
 this.city = city;
 this.state = state;
 this.zip = zip;
 this.phoneNumber = phoneNumber;
 }

 public int getStreetNum() {
 return streetNum;
 }

 public void setStreetNum(int streetNum) {

```



```

 this.streetNum = streetNum;
 }

 public java.lang.String getStreetName() {
 return streetName;
 }

 public void setStreetName(java.lang.String streetName) {
 this.streetName = streetName;
 }

 public java.lang.String getCity() {
 return city;
 }

 public void setCity(java.lang.String city) {
 this.city = city;
 }

 public java.lang.String getState() {
 return state;
 }

 public void setState(java.lang.String state) {
 this.state = state;
 }

 public int getZip() {
 return zip;
 }

 public void setZip(int zip) {
 this.zip = zip;
 }

 public WSIFPhone getPhoneNumber() {
 return phoneNumber;
 }

 public void setPhoneNumber(WSIFPhone phoneNumber) {
 this.phoneNumber = phoneNumber;
 }
}

```

### **WSIFPhone:**

```

public class WSIFPhone implements Serializable {

 //instance variables
 private int areaCode;
 private java.lang.String exchange;
 private java.lang.String number;

 //constructors
 public WSIFPhone () { }

 public WSIFPhone (int areaCode,
 java.lang.String exchange,
 java.lang.String number) {
 this.areaCode = areaCode;
 this.exchange = exchange;
 this.number = number;
 }

 public int getAreaCode() {
 return areaCode;
 }
}

```

```

 public void setAreaCode(int areaCode) {
 this.areaCode = areaCode;
 }

 public java.lang.String getExchange() {
 return exchange;
 }

 public void setExchange(java.lang.String exchange) {
 this.exchange = exchange;
 }

 public java.lang.String getNumber() {
 return number;
 }

 public void setNumber(java.lang.String number) {
 this.number = number;
 }
}

```

### **WSIFAddressBook:**

```

public class WSIFAddressBook {
 private Hashtable name2AddressTable = new Hashtable();

 public WSIFAddressBook() {
 }

 public void addEntry(String name, WSIFAddress address)
 {
 name2AddressTable.put(name, address);
 }

 public void addEntry(String firstName, String lastName, WSIFAddress address)
 {
 name2AddressTable.put(firstName+" "+lastName, address);
 }

 public WSIFAddress getAddressFromName(String name)
 throws IllegalArgumentException
 {
 if (name == null)
 {
 throw new IllegalArgumentException("The name argument must not be " +
 "null.");
 }
 return (WSIFAddress)name2AddressTable.get(name);
 }
}

```

The following code is the corresponding WSDL file for the Web service:

```

<?xml version="1.0" ?>

<definitions targetNamespace="http://www.ibm.com/namespace/wsif/samples/ab"
 xmlns:tns="http://www.ibm.com/namespace/wsif/samples/ab"
 xmlns:typens="http://www.ibm.com/namespace/wsif/samples/ab/types"
 xmlns:xsd="http://www.w3.org/2001/XMLSchema"
 xmlns:soap="http://schemas.xmlsoap.org/wsdl/soap/"
 xmlns:format="http://schemas.xmlsoap.org/wsdl/formatbinding/"
 xmlns:java="http://schemas.xmlsoap.org/wsdl/java/"
 xmlns:ejb="http://schemas.xmlsoap.org/wsdl/ejb/"
 xmlns="http://schemas.xmlsoap.org/wsdl/">

```

```

<types>
 <xsd:schema
 targetNamespace="http://www.ibm.com/namespace/wsif/samples/ab/types"
 xmlns:xsd="http://www.w3.org/2001/XMLSchema">

 <xsd:complexType name="phone">
 <xsd:element name="areaCode" type="xsd:int"/>
 <xsd:element name="exchange" type="xsd:string"/>
 <xsd:element name="number" type="xsd:string"/>
 </xsd:complexType>

 <xsd:complexType name="address">
 <xsd:element name="streetNum" type="xsd:int"/>
 <xsd:element name="streetName" type="xsd:string"/>
 <xsd:element name="city" type="xsd:string"/>
 <xsd:element name="state" type="xsd:string"/>
 <xsd:element name="zip" type="xsd:int"/>
 <xsd:element name="phoneNumber" type="typens:phone"/>
 </xsd:complexType>

 </xsd:schema>
</types>

<message name="AddEntryWholeNameRequestMessage">
 <part name="name" type="xsd:string"/>
 <part name="address" type="typens:address"/>
</message>

<message name="AddEntryFirstAndLastNamesRequestMessage">
 <part name="firstName" type="xsd:string"/>
 <part name="lastName" type="xsd:string"/>
 <part name="address" type="typens:address"/>
</message>

<message name="GetAddressFromNameRequestMessage">
 <part name="name" type="xsd:string"/>
</message>

<message name="GetAddressFromNameResponseMessage">
 <part name="address" type="typens:address"/>
</message>

<portType name="AddressBookPT">
 <operation name="addEntry">
 <input name="AddEntryWholeNameRequest"
 message="tns:AddEntryWholeNameRequestMessage"/>
 </operation>
 <operation name="addEntry">
 <input name="AddEntryFirstAndLastNamesRequest"
 message="tns:AddEntryFirstAndLastNamesRequestMessage"/>
 </operation>
 <operation name="getAddressFromName">
 <input name="GetAddressFromNameRequest"
 message="tns:GetAddressFromNameRequestMessage"/>
 <output name="GetAddressFromNameResponse"
 message="tns:GetAddressFromNameResponseMessage"/>
 </operation>
</portType>

<binding name="SOAPHttpBinding" type="tns:AddressBookPT">
 <soap:binding style="rpc"
 transport="http://schemas.xmlsoap.org/soap/http"/>
 <operation name="addEntry">
 <soap:operation soapAction=""/>
 <input name="AddEntryWholeNameRequest">
 <soap:body use="encoded"
 namespace="http://www.ibm.com/namespace/wsif/samples/ab"

```

```

 encodingStyle="http://schemas.xmlsoap.org/soap/encoding/" />
 </input>
</operation>
<operation name="addEntry">
 <soap:operation soapAction="" />
 <input name="AddEntryFirstAndLastNamesRequest">
 <soap:body use="encoded"
 namespace="http://www.ibm.com/namespace/wsif/samples/ab"
 encodingStyle="http://schemas.xmlsoap.org/soap/encoding/" />
 </input>
</operation>
<operation name="getAddressFromName">
 <soap:operation soapAction="" />
 <input name="GetAddressFromNameRequest">
 <soap:body use="encoded"
 namespace="http://www.ibm.com/namespace/wsif/samples/ab"
 encodingStyle="http://schemas.xmlsoap.org/soap/encoding/" />
 </input>
 <output name="GetAddressFromNameResponse">
 <soap:body use="encoded"
 namespace="http://www.ibm.com/namespace/wsif/samples/ab"
 encodingStyle="http://schemas.xmlsoap.org/soap/encoding/" />
 </output>
</operation>
</binding>

<binding name="JavaBinding" type="tns:AddressBookPT">
 <java:binding />
 <format:typeMapping encoding="Java" style="Java">
 <format:typeMap typeName="typens:address"
 formatType="com.ibm.www.namespace.wsif.samples.ab.types.WSIFAddress" />
 <format:typeMap typeName="xsd:string" formatType="java.lang.String" />
 </format:typeMapping>
 <operation name="addEntry">
 <java:operation
 methodName="addEntry"
 parameterOrder="name address"
 methodType="instance" />
 <input name="AddEntryWholeNameRequest" />
 </operation>
 <operation name="addEntry">
 <java:operation
 methodName="addEntry"
 parameterOrder="firstName lastName address"
 methodType="instance" />
 <input name="AddEntryFirstAndLastNamesRequest" />
 </operation>
 <operation name="getAddressFromName">
 <java:operation
 methodName="getAddressFromName"
 parameterOrder="name"
 methodType="instance"
 returnPart="address" />
 <input name="GetAddressFromNameRequest" />
 <output name="GetAddressFromNameResponse" />
 </operation>
</binding>

<binding name="EJBBinding" type="tns:AddressBookPT">
 <ejb:binding />
 <format:typeMapping encoding="Java" style="Java">
 <format:typeMap typeName="typens:address"
 formatType="com.ibm.www.namespace.wsif.samples.ab.types.WSIFAddress" />
 <format:typeMap typeName="xsd:string" formatType="java.lang.String" />
 </format:typeMapping>
 <operation name="addEntry">
 <ejb:operation

```

```

 methodName="addEntry"
 parameterOrder="name address"
 interface="remote"/>
 <input name="AddEntryWholeNameRequest"/>
</operation>
<operation name="addEntry">
 <ejb:operation
 methodName="addEntry"
 parameterOrder="firstName lastName address"
 interface="remote"/>
 <input name="AddEntryFirstAndLastNamesRequest"/>
</operation>
<operation name="getAddressFromName">
 <ejb:operation
 methodName="getAddressFromName"
 parameterOrder="name"
 interface="remote"
 returnPart="address"/>
 <input name="GetAddressFromNameRequest"/>
 <output name="GetAddressFromNameResponse"/>
</operation>
</binding>
<service name="AddressBookService">
 <port name="SOAPPort" binding="tns:SOAPHttpBinding">
 <soap:address
 location="http://localhost/wsif/samples/addressbook/soap/servlet/rpcrouter"/>
 </port>
 <port name="JavaPort" binding="tns:JavaBinding">
 <java:address className="services.addressbook.WSIFAddressBook"/>
 </port>
 <port name="EJBPort" binding="tns:EJBBinding">
 <ejb:address className="services.addressbook.ejb.AddressBookHome"
 jndiName="ejb/samples/wsif/AddressBook"
 classLoader="services.addressbook.ejb.AddressBook.ClassLoader"/>
 </port>
</service>
</definitions>

```

## Using complex types

WSIF supports user-defined complex types through the mapping of complex types to Java classes.

You specify this mapping manually or automatically as described in the following sections:

- Manual mapping of complex types.
- Automatic mapping of complex types.

Any calls to the WSIFService mapType and mapPackage methods used for manual mapping override any equivalent mapping information that is produced automatically. This override helps to maintain backwards compatibility, and also accommodates less standard mappings.

### Manual mapping of complex types

The method to use when you create these mappings manually depends on the provider that is used. For the Java and EJB providers, the mappings are specified in the WSDL file in the binding element. The following example provides the syntax for specifying the mapping:

```

<binding >
 <ejb:binding|java:binding/>
 <format:typeMapping style="Java" encoding="Java"/>?
 <format:typeMap name="qname" formatType="nmtoken"/>*
 </format:typeMapping>
 ...
</binding>

```

In this example:

- A question mark (“?”) means “optional” and an asterisk (“\*”) means “0 or more”.
- The `format:typeMap` **name** attribute is a qualified name of a complex type or simple type used by one of the operations.
- The `format:typeMap` **formatType** attribute is the fully qualified class name for the Java class to which the element specified by **name** maps.

If you use the Apache SOAP provider then you specify the mapping of a complex type to a Java class in the client code through two methods on the `org.apache.wsif.WSIFService` interface:

```
public void mapType(QName elementType, Class javaType)
```

and

```
public void mapPackage(String namespaceURI, String packageName)
```

Use the **mapType** method to specify a mapping between an XML schema element and a Java class. The method takes a `QName` representing the complex type or simple type, and the corresponding Java class to which it maps.

Use the **mapPackage** method to specify a more general mapping between a namespace and a Java package. Any custom, complex or simple type whose namespace matches that of the mapping is mapped to a Java class in the corresponding package. The name of the actual class is derived from the name of the complex type using standard XML to Java naming conventions.

### Automatic mapping of complex types

For complex types defined in the WSDL, where a generated bean is used to represent this type in Java, the Web Services Invocation Framework (WSIF) programming model requires that a call is made to the `WSIFService.mapType()` method. This call tells WSIF the package and class name of the bean representing the XML schema type that is identified with a `QName`. To make things easier, the `WSIFService.mapPackage()` method provides a mechanism to specify a wildcard version of this, where any class within a specified package is mapped to the namespace of a `QName`. This is a mechanism for manually mapping an XML schema type to a Java class and back again (one mapping entry provides a bidirectional mapping).

There are many ways to convert a `QName` representing an XML schema type name to a Java package name and class. To enable automatic type mapping, set the `WSIF_FEATURE_AUTO_MAP_TYPES` feature on the `WSIFServiceFactory` instance:

```
WSIFServiceFactory factory = WSIFServiceFactory.newInstance();
factory.setFeature(WSIFConstants.WSIF_FEATURE_AUTO_MAP_TYPES, new Boolean(true));
```

WSIF maps types by converting the URI part of the XML schema type `<tt>QName</tt>` to a package name, and converting the local part to a class name. WSIF does this mapping using the `WSIFUtils` methods `<tt>getPackageNameFromNamespaceURI</tt>` and `<tt>getJavaClassNameFromXMLName</tt>`.

## Using the Java Naming and Directory Interface (JNDI)

This example task shows you how to use WSIF to bind a reference to a Web service, then look up the reference using JNDI.

You access a Web service through information provided in the WSDL document for the service. If you do not know where to find the WSDL document for the service, but you know that it has been registered in a UDDI registry, then you look it up in the registry. Java programs access Java objects and resources in a similar manner, but using a JNDI interface.

The following example shows how, using the Web Services Invocation Framework (WSIF), you can bind a reference to a Web service then look up the reference using JNDI.

### Specifying the argument values for the Web service

The Web service is represented in WSIF by an instance of the `org.apache.wsif.naming.WSIFServiceRef` class. This simple Referenceable object has the following constructor:

```
public WSIFServiceRef(
 String WSDL,
 String sNS,
 String sName,
 String ptNS,
 String ptName)
{
 [...]
}
```

In this example

- *WSDL* is the location of the WSDL file containing the definition of the service.
- *sNS* is the full namespace for the service definition (you can specify `null` if only one service is defined in the WSDL file).
- *sName* is the local name for the service definition (you can specify `null` if only one service is defined in the WSDL file).
- *ptNS* is the full namespace for the port type within the service that you want to use (you can specify `null` if only one port type is available for the service).
- *ptName* is the local name for the port type (you can specify `null` if only one port type is available for the service).

For example, if the WSDL file for the Web service is available from the Web address `http://localhost/WSDL/Example.WSDL` and contains the following service and port type definitions:

```
<definitions targetNamespace="http://hostname/namespace/example"
 xmlns:abc="http://hostname/namespace/abc"
[...]
 <portType name="ExamplePT">
 <operation name="exampleOp">
 <input name="exampleInput" message="tns:ExampleInputMsg"/>
 </operation>
 </portType>
[...]
 <service name="abc:ExampleService">
[...]
 </service>
[...]
</definitions>
```

You can specify the following argument values for the `WSIFServiceRef` class:

- *WSDL* is `http://localhost/WSDL/Example.WSDL`
- *sNS* is `http://hostname/namespace/abc`
- *sName* is `ExampleService`
- *ptNS* is `http://hostname/namespace/example`
- *ptName* is `ExamplePT`

### Binding the service using JNDI

To bind the service reference in the naming directory using JNDI, you can use the `com.ibm.websphere.naming.JndiHelper` class in WebSphere Application Server:

```
[...]
import com.ibm.websphere.naming.JndiHelper;
import org.apache.wsif.naming.*;
```



```
[...]
try {
 Context startingContext = new InitialContext();
 WSIFServiceRef ref = new WSIFServiceRef("http://localhost/WSDL/Example.WSDL",
 "http://hostname/namespace/abc"
 "ExampleService",
 "http://hostname/namespace/example",
 "ExamplePT");

 JndiHelper.recursiveRebind(startingContext,
 "myContext/mySubContext/myServiceRef", ref);

}
catch (NamingException e) {
 // Handle error.
}
[...]
```

## Looking up the service using JNDI

The following code fragment shows the lookup of a service using JNDI:

```
[...]
try {
[...]
```

```
 InitialContext ic = new InitialContext();
 WSIFService myService =
 (WSIFService) ic.lookup("myContext/mySubContext/myServiceRef");
[...]
```

```
}
catch (NamingException e) {
 // Handle error.
}
[...]
```

## Passing SOAP messages with attachments using WSIF

The W3C SOAP Messages with Attachments document describes a standard way to associate a SOAP message with one or more attachments in their native format (for example GIF or JPEG) by using a multipart MIME structure for transport. It defines specific use of the "Multipart/Related" MIME media type, and rules for the use of URI references to entities bundled within the MIME package. It thereby outlines a technique for carrying a SOAP 1.1 message within a MIME multipart/related message in such a way that the SOAP processing rules for a standard SOAP message are not changed.

The Web Services Invocation Framework (WSIF) supports passing attachments in a MIME message using the SOAP provider. The attachment is a `javax.activation.DataHandler` object. The `mime:multipartRelated`, `mime:part` and `mime:content` tags are used to describe the attachment in the WSDL.

For more information, see the following topics:

- SOAP messages with attachments - Writing the WSDL extensions.
- SOAP messages with attachments - Passing attachments to WSI.
- SOAP messages with attachments - Working with types and type mappings.

The following scenarios are not supported:

- Using DIME.
- Passing in `javax.xml.transform.Source` and `javax.mail.internet.MimeMultipart`.
- Using the `mime:mimeXml` WSDL tag.
- Nesting a `mime:multipartRelated` tag inside a `mime:part` tag.
- Using types that extend `DataHandler`, `Image`, and so on.
- Using types that contain `DataHandler`, `Image`, and soon.
- Using Arrays or Vectors of `DataHandlers`, `Images`, and so on.
- Using multiple in/out or output attachments.

The MIME headers from the incoming message are not preserved for referenced attachments. The outgoing message contains new MIME headers for Content-Type, Content-Id and Content-Transfer-Encoding that are created by WSIF.

## SOAP messages with attachments - Writing the WSDL extensions

The following example WSDL illustrates a simple operation that has one attachment called attch:

```
<binding name="MyBinding" type="tns:abc" >
 <soap:binding style="rpc" transport="http://schemas.xmlsoap.org/soap/http"/>
 <operation name="MyOperation">
 <soap:operation soapAction=""/>
 <input>
 <mime:multipartRelated>
 <mime:part>
 <soap:body use="encoded" namespace="http://mynamespace"
 encodingStyle="http://schemas.xmlsoap.org/soap/encoding"/>
 </mime:part>
 <mime:part>
 <mime:content part="attch" type="text/html"/>
 </mime:part>
 </mime:multipartRelated>
 </input>
 </operation>
</binding>
```

In this type of WSDL extension:

- There must be a part attribute (in this example attch) on the input message for the operation (in this example MyOperation). There can be other input parts to MyOperation that are not attachments.
- In the binding input there must either be a <soap:body> tag or a <mime:multipartRelated> tag, but not both.
- For MIME messages, the <soap:body> tag is inside a <mime:part> tag. There must only be one <mime:part> tag that contains a <soap:body> tag in the binding input and that must not contain a <mime:content> tag as well, because a content type of text/xml is assumed for the <soap:body> tag.
- There can be multiple attachments in a MIME message, each described by a <mime:part> tag.
- Each <mime:part> tag that does not contain a <soap:body> tag contains a <mime:content> tag that describes the attachment itself. The type attribute inside the <mime:content> tag is not checked or used by the Web Services Invocation Framework (WSIF). It is there to suggest to the application using WSIF what the attachment contains. Multiple <mime:content> tags inside a single <mime:part> tag means that the backend service expects a single attachment with a type specified by one of the <mime:content> tags inside that <mime:part> tag.
- The parts="..." attribute (optional) inside the <soap:body> tag is assumed to contain the names of all the MIME parts as well as the names of all the SOAP parts in the message.

## SOAP messages with attachments - Passing attachments to WSIF

The following code fragment can invoke the service described by the example WSDL in the topic writing the WSDL extensions:

```
import javax.activation.DataHandler;
. . .
DataHandler dh = new DataHandler(new FileDataSource("myimage.jpg"));
WSIFServiceFactory factory = WSIFServiceFactory.newInstance();
WSIFService service = factory.getService("my.wsdl",null,null,"http://mynamespace","abc");
WSIFOperation op = service.getPort().createOperation("MyOperation");
WSIFMessage in = op.createInputMessage();
in.setObjectPart("attch",dh);
op.executeInputOnlyOperation(in);
```

The associated type mapping in the DeploymentDescriptor.xml file depends upon your SOAP server. For example if you use Tomcat with SOAP 2.3, then the DeploymentDescriptor.xml file contains the following type mapping:

```

<isd:mappings>
<isd:map encodingStyle="http://schemas.xmlsoap.org/soap/encoding/"
 xmlns:x="http://mynamespace"
 qname="x:datahandler"
 javaType="javax.activation.DataHandler"
 java2XMLClassName="org.apache.soap.encoding.soapenc.MimePartSerializer"
 xml2JavaClassName="org.apache.soap.encoding.soapenc.MimePartSerializer" />
</isd:mappings>

```

In this case, the backend service is invoked with the following signature:

```
public void MyOperation(DataHandler dh);
```

You can also use stubs to pass attachments into the Web Services Invocation Framework (WSIF):

```

DataHandler dh = new DataHandler(new FileDataSource("myimage.jpg"));
WSIFServiceFactory factory = WSIFServiceFactory.newInstance();
WSIFService service = factory.getService("my.wsdl", null, null, "http://mynamespace", "abc");
MyInterface stub = (MyInterface)service.getStub(MyInterface.class);
stub.MyOperation(dh);

```

Attachments can also be returned from an operation, but at present only one attachment can be returned as the return parameter.

## SOAP messages with attachments - Working with types and type mappings

By default, attachments are passed into the Web Services Invocation Framework (WSIF) as `DataHandler` objects. If the part on the message that is the `DataHandler` object maps to a `<mime:part>` tag in the WSDL, then WSIF automatically maps the fully qualified name of the WSDL type to the `DataHandler` class and sets up that type mapping with the SOAP provider.

In your WSDL, you might have defined a schema for the attachment (for instance as a `binary[]` type). WSIF silently ignores this mapping and treats the attachment as a `DataHandler` object, unless you explicitly issue a `mapType()` method. WSIF lets the SOAP provider set the MIME content type based on the type of the `DataHandler` object, instead of the `type` attribute specified for the `<mime:content>` tag in the WSDL.

## Interacting with the J2EE container in WebSphere Application Server

Interaction with a container is limited to the following aspects:

- Using the application server administrative console to define Web services to WebSphere Application Server. This task is described in *Using the Java Naming and Directory Interface (JNDI) and WSIF system management and administration*. As part of the definition of a service, the administrator might define a “preferred port”.
- Using the Web Services Invocation Framework (WSIF) to make log and trace calls to the J2EE services in WebSphere Application Server, as described in *Trace and logging for WSIF*.
- Using WSIF providers to access Java 2 platform, Enterprise Edition (J2EE) services. For example using the EJB provider to access the Java Naming and Directory Interface (JNDI) and make calls to remote enterprise beans.
- Using WSIF to wrap the use of container services so that, when WSIF is run in an unmanaged (thin) environment, the operation can succeed.

## Running WSIF as a client

The Web Services Invocation Framework (WSIF) runs in the WebSphere Application Server application client container, and in similar clients from other suppliers.

To simplify the process of launching client applications in the WebSphere Application Server application client, use the `launchClient` tool as described in *Running application clients*.

---

## WSIF system management and administration

The Web Services Invocation Framework (WSIF) is provided as a stand-alone JAR file named `wsif.jar`. The JAR file contains the core WSIF classes, and the Java, EJB, SOAP over HTTP and SOAP over JMS providers. Additional providers are packaged as separate JAR files.

When you install WebSphere Application Server, the `wsif.jar` file is put on the WebSphere or Java Virtual Machine (JVM) class path.

WSIF requires no further configuration. WSIF is a thin abstraction layer between application code and the relevant invocation infrastructure.

For specific information on other aspects of managing your WSIF system, see the following topics:

- Maintaining the WSIF properties file
- Enabling security for WSIF
- Trace and logging for WSIF
- Troubleshooting the Web Services Invocation Framework
- WSIF (Web Services Invocation Framework) messages

### Maintaining the WSIF properties file

The Web Services Invocation Framework (WSIF) properties are stored in the `wsif.jar` file, in a properties file named `wsif.properties`. This properties file is kept on the class path, so that WSIF can find it and the client administrator can use it to configure WSIF.

Here is a copy of the initial contents of the `wsif.properties` file. All the possible properties are listed and described.

```
Two properties are used to override which WSIFProvider is selected when there
exists multiple providers supporting the same namespace URI. These properties are:
#
wsif.provider.default.CLASSNAME=N
wsif.provider.uri.M.CLASSNAME=URI
#
CLASSNAME is the WSIFProvider class name
N is the number of following default wsif.provider.uri.M.CLASSNAME properties
M is a number from 1 to N to uniquely identify each wsif.provider.uri.M.CLASSNAME
property key.
For example the following two properties would override the default SOAP provider
to be the Apache SOAP provider:
#
wsif.provider.default.org.apache.wsif.providers.soap.ApacheSOAP.WSIFDynamicProvider_ApacheSOAP=1
wsif.provider.uri.1.org.apache.wsif.providers.soap.ApacheSOAP.WSIFDynamicProvider_ApacheSOAP=\
http://schemas.xmlsoap.org/wsdl/soap/
#

maximum number of milliseconds to wait for a response to a synchronous request.
Default value if not defined is to wait forever.
wsif.syncrequest.timeout=10000

maximum number of seconds to wait for a response to an async request.
if not defined on invalid defaults to no timeout
wsif.asyncrequest.timeout=60
```

### Enabling security for WSIF

The Web Services Invocation Framework (WSIF) interacts with a security manager in the following ways:

- WSIF runs in the Java 2 platform, Enterprise Edition (J2EE) security context without modification.
- When WSIF is run under a J2EE container, port implementations can use the security context to pass on security tokens or credentials as necessary.

- WSIF implementations can automatically convert J2EE security context into appropriate context for onward services.

For WSIF to interact effectively with the WebSphere Application Server security manager, enable the following permission in the `was.policy` file: **FilePermission** to load the WSDL. This permission is required when a WSDL file is referred to using the `file://` protocol.

## Troubleshooting the Web Services Invocation Framework

For information on resolving WebSphere-level problems, see "Diagnosing and fixing problems" in the Information Center.

To identify and resolve Web Services Invocation Framework (WSIF)-related problems, you can use the standard WebSphere Application Server trace and logging facilities. If you encounter a problem that you think might be related to WSIF, you can check for error messages in the WebSphere Application Server administrative console, and in the application server `stdout.log` file. You can also enable the application server debug trace to provide a detailed exception dump.

A list of the WSIF run-time system messages, with details of what each message means, is provided in Message reference for WSIF.

Here is a checklist of major WSIF activities, with advice on common problems associated with each activity:

### Create service

Handcrafted WSDL can cause numerous problems. To help ensure that your WSDL is valid, use a tool such as WebSphere Studio to create your service.

### Define transport mechanism

For the Java Messaging Service (JMS), check that you have set up the Java Naming and Directory Interface (JNDI) correctly, and created the necessary connection factories and queues.

For SOAP, make sure that the deployment descriptor file `dds.xml` is correct - preferably by creating it using WebSphere Studio or similar tooling.

### Create client - Java code

Follow the correct format for creating a WSIF service, port, operation and message. For examples of correct code, see the Address Book Sample.

### Compile code (client and service)

Check that the build path against code is correct, and that it contains the correct levels of JAR files.

Create a valid EAR file for your service in preparation for deployment to a Web server.

### Deploy service

When you install and deploy the service EAR file, check carefully any messages given when the service is deployed.

### Server setup and start

Make sure that the WebSphere Application Server `server.policy` file (in the `/properties` directory) has the correct security settings. For more information, see Enabling security for WSIF.

### WSIF setup

Check that the `wsif.properties` file is correctly set up. For more information, see Maintaining the WSIF properties file.

### Run client

Either check that you have defined the class path correctly to include references to your client classes, WSIF JAR files and any other necessary JAR files, or (preferably) run your client using the WebSphere Application Server `launchClient` tool.

Here is a list of common errors, and information on their probable causes:

- **"No class definition" errors received when running client code.**

This problem usually indicates an error in the class path setup. Check that the relevant JAR files are included.

- **“Cannot find WSDL” error.**

Some likely causes are:

- The application server is not running.
- The server location and port number in the WSDL are not correct.
- The WSDL is badly formed (check the error messages in the application server `stdout.log` file).
- The application server has not been restarted since the service was installed.

You might also try the following checks:

- Can you load the WSDL into your Web browser from the location specified in the error message?
- Can you load the corresponding WSDL binding files into your Web browser?

- **Your Web service EAR file does not install correctly onto the application server.**

It is likely that the EAR file is badly formed. Verify the installation by completing the following steps:

- For an EJB binding, run the WebSphere Application Server tool `\bin\dumpnamespace`. This tool lists the current contents of the JNDI directory.
- For a SOAP over HTTP binding, open the `http://pathToServer/WebServiceName/admin/list.jsp` page (if you have the SOAP administration pages installed). This page lists all currently installed Web services.
- For a SOAP over JMS binding, complete the following checks:
  - Check that the queue manager is running.
  - Check that the necessary queues are defined.
  - Check the JNDI setup.
  - Use the “display context” option in the `jmsadmin` tool to list the current JNDI definitions.
  - Check that the Remote Procedure Call (RPC) router is running.

- **There is a permissions problem or security error.**

Check that the WebSphere Application Server `server.policy` file (in the `/properties` directory) has the correct security settings. For more information, see *Enabling security for WSIF*.

- **Using WSIF with multiple clients causes a SOAP parsing error.**

Before you deploy a Web service to WebSphere Application Server, you must decide on the scope of the Web service. The deployment descriptor file `dds.xml` for the Web service includes the following line:

```
<isd:provider type="java" scope="Application"
```

You can set the `Scope` attribute to `Application` or `Session`. The default setting is `Application`, and this value is correct if each request to the Web service does not require objects to be maintained for longer than a single instance. If `Scope` is set to `Application` the objects are not available to another request during the execution of the single instance, and they are released on completion. If your Web service needs objects to be maintained for multiple requests, and to be unique within each request, you must set the scope to `Session`. If `Scope` is set to `Session`, the objects are not available to another request during the life of the session, and they are released on completion of the session. If scope is set to `Application` instead of `Session`, you might get the following SOAP error:

```
SOAPException: SOAP-ENV:ClientParsing error, response was:
FWK005 parse may not be called while parsing.;
nested exception is:
```

```
[SOAPException: faultCode=SOAP-ENV:Client; msg=Parsing error, response was:
```

```
FWK005 parse may not be called while parsing.;
targetException=org.xml.sax.SAXException:
FWK005 parse may not be called while parsing.]
```

- **5.1+ Using the same names for JMS messaging queues and queue connection factories that run on application servers on different machines can cause JNDI lookup errors.** You should not use the same names for messaging queues and queue connection factories that run on application servers on different machines, because WSIF always looks first for JMS destinations locally, and only uses the full JNDI reference if it cannot find the destination locally. For example, if you run a Web service on a remote machine, and have an application server running locally that uses the same names



for the messaging queues and queue connection factories, then WSIF will find and use the local queues even if the remote JNDI destination is provided in full in the WSDL service definition.

## Trace and logging for WSIF

If you want to enable trace for the Web Services Invocation Framework (WSIF) API within WebSphere Application Server, and have trace, stdout and stderr for the application server written to a well-known location, see "Enabling trace" in the Information Center.

WSIF offers trace points at the opening and closing of ports, the invocation of services, and the responses from services.

To trace the WSIF API, you need to specify the following trace string:

```
wsif=all=enabled
```

WSIF also includes a SimpleLog utility through which you can run trace when using WSIF outside of WebSphere Application Server. To enable this utility, complete the following steps:

1. Create a file named `commons-logging.properties` with the following contents:

```
org.apache.commons.logging.LogFactory=org.apache.commons.logging.impl.LogFactoryImpl
org.apache.commons.logging.Log=org.apache.commons.logging.impl.SimpleLog
```

2. Create a file named `simplelog.properties` with the following contents:

```
org.apache.commons.logging.simplelog.defaultlog=trace
org.apache.commons.logging.simplelog.showShortLogname=true
org.apache.commons.logging.simplelog.showdatetime=true
```

3. Put both these files, and the `commons-logging.jar` file, on the class path.

The SimpleLog utility writes trace to the `System.err` file.

## WSIF (Web Services Invocation Framework) messages

This topic contains a list of the WSIF run-time system messages, with details of what each message means.

WebSphere system messages are logged from a variety of sources, including application server components and applications. Messages logged by application server components and associated IBM products start with a unique message identifier that indicates the component or application that issued the message.

For more information about the message identifier format, see the topic "Message reference" in the Information Center.

### **WSIF0001E: An extension registry was not found for the element type "{0}"**

**Explanation:** Parameters: {0} element type. No extension registry was found for the element type specified.

**User Response:** Add the appropriate extension registry to the port factory in your code.

### **WSIF0002E: A failure occurred in loading WSDL from "{0}"**

**Explanation:** Parameters: {0} location of the WSDL file. The WSDL file could not be found at the location specified or did not parse correctly

**User Response:** Check that the location of the WSDL file is correct. Check that any network connections required are available. Check that the WSDL file contains valid WSDL.

### **WSIF0003W: An error occurred finding pluggable providers: {0}**

**Explanation:** Parameters: {0} specific details about the error. There was a problem locating a WSIF pluggable provider using the J2SE 1.3 JAR file extensions to support service providers architecture. The WSIF trace file will contain the full exception details.



**User Response:** Verify that a META-INF/services/org.apache.wsif.spi.WSIFProvider file exists in a provider jar, that each class referenced in the META-INF file exists in the class path, and that each class implements org.apache.wsif.spi.WSIFProvider. The class in error will be ignored and WSIF will continue locating other pluggable providers.

**WSIF0004E: WSDL contains an operation type “{0}” which is not supported for “{1}”**

**Explanation:** Parameters: {0} name of the operation type specified. {1} name of the portType for the operation. An operation type which is not supported has been specified in the WSDL.

**User Response:** Remove any operations of the unsupported type from the WSDL. If the operation is required then make sure all messages have been correctly specified for the operation.

**WSIF0005E: An error occurred when invoking the method “{1}” . (“{0}” )**

**Explanation:** Parameters: {0} name of communication type. For example EJB or Apache SOAP. {1} name of the method that failed. An error was encountered when invoking a method on the Web service using the communication shown in brackets.

**User Response:** Check that the method exists on the Web service and that the correct parts have been added to the operation as described in the WSDL. Network problems might be a cause if the method is remote and so check any required connections.

**WSIF0006W: Multiple WSIFProvider found supporting the same namespace URI “{0}” . Found (“{1}” )**

**Explanation:** Parameters: {0} the namespace URI. {1} a list of the WSIFProvider found.. There are multiple org.apache.wsif.spi.WSIFProvider classes in the service provider path that support the same namespace URI.

**User Response:** A following WSIF0007I message will be issued notifying which WSIFProvider will be used. Which WSIFProvider is chosen is based on settings in the wsif.properties file, or if not defined in the properties, the last WSIFProvider found will be used. See the wsif.properties file for more details on how to define which provider should be used to support a namespace URI.

**WSIF0007I: Using WSIFProvider “{0}” for namespaceURI “{1}”**

**Explanation:** Parameters: {0} the classname of the WSIFProvider being used. {1} the namespaceURI the provider will be used to support.. Either a previous WSIF0006W message has been issued or the SetDynamicWSIFProvider method has been used to override the provider used to support a namespaceURI.

**User Response:** None. See also WSIF0006W.

**WSIF0008W: WSIFDefaultCorrelationService removing correlator due to timeout. ID:“{0}”**

**Explanation:** Parameters: {0} the ID of the correlator being removed from the correlation service. A stored correlator is being removed from the correlation service due to its timeout expiring.

**User Response:** Determine why no response has been received for the asynchronous request within the timeout period. The wsif.asyncrequest.timeout property of the wsif.properties file defines the length of the timeout period.

**WSIF0009I: Using correlation service - “{0}”**

**Explanation:** Parameters: {0} the name of the correlation service being used. This identifies the name of the correlation service that will be used to process asynchronous requests.

**User Response:** None. If a correlation service other than the default WSIF supplied one is required, ensure that it is correctly registered in the JNDI java:comp/wsif/WSIFCorrelationService namespace.

**WSIF0010E: Exception thrown while processing asynchronous response - “{0}”**

**Explanation:** Parameters: {0} the error message string of the exception. While processing the response from an executeRequestResponseAsync call an exception was thrown.

**User Response:** Use the exception error message string to determine the cause of the error. The WSIF trace will have more details on the error including the exception stack trace.

**WSIF0011I: Preferred port “{0}” was not available**

**Explanation:** Parameters: {0} the user’s preferred port. The preferred port set by the user on org.apache.wsif.WSIFService is not available

**User Response:** None unless this message appears for long periods of time in which case the user might want to pick a different port as their preferred port.

---

## WSIF API

The WSIF API supports the invocation of services defined in WSDL. WSIF is intended for use in both WSIF clients and Web service intermediaries.

The WSIF API is driven by the abstract service description in WSDL; it is completely independent of the actual binding used. This independence makes the API more natural to work with because it uses WSDL terms to refer to message parts, operations, and so on.

The WSIF API was designed for the WSDL usage model: Pick a port that supports the port type needed, then invoke the operation by providing the necessary abstract input message consisting of the required parts, without worrying about how the message is mapped to a specific binding protocol.

Other Web service APIs, for example SOAP APIs, are not designed on WSDL, but for a specific binding protocol with its associated syntax; for example, target URIs and encoding styles.

The WSIF API main interfaces are described in the following topics:

- Creating a message for sending to a port (the `WSIFMessage` interface).
- WSIF API reference: Finding a port factory or service (the `WSIFService` interface and the `WSIFServiceFactory` class).
- WSIF API reference: Using ports (the `WSIFPort` interface and the `WSIFOperation` interface).

**Note:** You must ensure that your application uses only one thread to call WSIF.

For additional technical details of the WSIF API, see the WSIF Javadoc.

### WSIF API reference: Creating a message for sending to a port

For message management (that is, message construction and parsing) the underlying API is modeled on WSDL semantics. There is a simple and direct mapping from the WSDL model to the Web Services Invocation Framework (WSIF) classes.

In WSDL, a message describes the abstract type of the input or output to an operation. The corresponding WSIF class is `WSIFMessage`, which represents in memory the actual input or output of an operation. A `WSIFMessage` class is a container for a set of named parts. The `WSIFMessage` interface separates the actual representation of the data from the abstract type defined by WSDL. WSDL defines messages as XML schema types. There are two natural ways to represent a WSDL message in a run-time environment:

- The generated Java class, based on a WSDL to Java mapping such as that provided by a Java API for XML-based remote procedure call (JAX-RPC).
- The XML representation of the data, for example using SOAP Encoding.

Each option offers benefits in different scenarios. The Java class is the natural approach when WSIF is used in a standard Java client. However, in other scenarios where WSIF is used in an intermediary, it might be more efficient to keep a WSDL message in the SOAP encoded format.

The style used to define messages must be consistent within the message, so all the parts in one message must be consistent. A string - `getRepresentationStyle()` - always returns `null`. This indicates that parts on this `WSIFMessage` class are represented as Java objects.

You add parts to a `WSIFMessage` class with the `setObjectPart` or `setTypePart` methods. Each part is named. Part names within a message are unique. If you set a part more than once, the last setting is the one that is used.

You retrieve parts by name from a `WSIFMessage` class with the `getObjectPart` or `getTypePart` methods. If the named part does not exist, the method returns a `WSIFException` exception.

You can use Iterators to retrieve parts from the message through the `getParts()` and `getPartNames()` methods.

The order in which you set the parts is not important, but the message implementation might be more efficient if the parts are set in the parameter order specified by WSDL.

WSIFMessage classes are cloneable and serializable. If the parts set are not cloneable, the implementation can try to clone them using serialization. If the parts are not serializable either, then a `CloneNotSupportedException` exception is thrown if cloning is attempted.

WSIFMessage classes can be sent between Java Virtual Machines (JVMs).

In addition to the containing parts, a WSIFMessage class also has a message name. This is required for operation overloading, which is supported by WSDL and WSIF.

Here is the Javadoc for the WSIFMessage interface.

## WSIF API reference: Finding a port factory or service

To find a port you use the WSIFService interface, which is a factory for ports.

The port factory models and supports the WSDL approach in which a service is available on one or more ports. The factory hides the implementation of the port from the user. The Web Services Invocation Framework (WSIF) supports dynamic ports that are based on a particular protocol and transport, and configured using the WSDL at run-time. For example, the dynamic SOAP port can invoke any SOAP service based on the WSDL description of that service. Using this service you can hide and modify the set of available ports at run-time.

Here is the WSIFService interface.

To find a service from a WSDL document at a Web address, or from a code-generated code base, you can use the WSIFServiceFactory class.

### WSIFService interface

The WSIFService interface is responsible for generating an instance of the WSIFOperation interface to use for a particular invocation of a service operation.

The Web Services Invocation Framework (WSIF) service stores a list of providers that can each generate a WSIF operation for a particular WSDL binding. This service looks up providers by the provider type. For example the service knows about one provider that handles SOAP ports and other providers that handle Java ports that you define. In a managed environment, the container can configure the WSIFService interface.

Here is the Javadoc for the WSIFService interface.

A WSIFService implementation can choose a preferred port based on a number of criteria. The WSIFService implementation can set the preferred port, or it can be set by calling the `setPreferredPort` method.

The `getPort` method returns an instance of the WSIFPort class that is used to invoke a service on the port. Variants of the `getPort` method are used to define the characteristics of the port to be created:

- the `getPort` method with no arguments returns the preferred port.
- the `getPort` method with a string argument returns the port named by the string containing the WSDL identifier for the selected port.

The return value is `null` if the port name is not valid.

If a port is chosen (either by the `WSIFService` implementation, or by the `setPreferredPort` method), then the `WSIFService` implementation validates that the relevant provider exists and is configured. If the provider fails this validation check, the `WSIFService` interface chooses any other port for which a provider is defined. For example, if the preferred port is SOAP over JMS but the JMS libraries are not available, then WSIF chooses another port. If no preferred port is set, or the preferred port is not available, the `WSIF` implementation chooses the first available port listed in the WSDL.

The `getAvailablePortNames()` method returns, as an iteration of strings, the list of WSDL port names filtered by the set of available providers.

The `getDefinition()` method returns the WSDL definition for the service. If the WSDL definition is not available, this method returns `null`.

### **WSIFServiceFactory class**

To find a service from a WSDL document at a Web address, or from a code-generated code base, you can use the `WSIFServiceFactory` class.

**Note:** When you create a `WSIFService` interface from a `WSIFServiceFactory` class, you can specify a `ClassLoader` object to use in locating the WSDL file. You need to specify this object when the WSDL file is in a JAR file. In such a case, specify the location of the WSDL file relative to the root of the JAR file, using forward slashes (/) with the preceding slash removed.

For example:

```
com/myCompany/wsd1/MyWSDLFile.wsdl
```

rather than

```
/com/myCompany/wsd1/MyWSDLFile.wsdl
```

Here is the Javadoc for the `WSIFServiceFactory` class.

The `WSIFServiceFactory` class returns `null` if no service is found with that identifier.

### **WSIF API reference: Using ports**

A `WSIFPort` interface handles the details of invoking an operation. The port provides access to the actual implementation of the service.

A WSDL document can provide many different WSDL bindings, and these bindings can drive multiple ports. The client can choose a port, the service stub can choose a port, or the Web Services Invocation Framework (WSIF) can choose a default port.

The port offers an interface to retrieve an `Operation` object. A `WSIFOperation` interface offers the ability to execute the given operation.

If the port is serialized and deserialized at a later time, then WSIF ensures that the client provides the correct information to the server to identify the instance. If the server instance is no longer available, then it is up to the server to decide whether to throw a fault or provide a new instance. That behavior can depend on the type of service.

For example, for an enterprise bean the `WSIFPort` interface stores the EJB Home, and uses it to select the bean before each invocation. It is the responsibility of the client to serialize or maintain the port instance if it wants instance support. The client must create a new operation and messages for each invocation.

Here is the `WSIFPort` interface.

Here is the WSIFOperation interface.

### **WSIFPort interface**

The port implements a factory method for the WSIFOperation interface.

Here is the Javadoc for the WSIFPort interface.

The createOperation(String) method returns a new instance of a WSIFOperation object. If the operationName value is not valid or the operation is overloaded, then the method throws an exception.

The createOperation(String, String, String) method supports overloaded WSDL operations. You can overload based on the input parameters, but not on the output parameters.

It is the duty of the client to call the close method when a port is no longer in use. In many cases, where the transport is sessionless, like HTTP, this has no effect. However, if the port is using a session-based protocol such as MQSeries, Java Messaging Service (JMS), or External Call Interface (ECI), this supports the port in caching an open connection to the server and then closing it as required. Responsibly-written applications will call the close method if appropriate.

### **WSIFOperation interface**

You use the WSIFOperation interface to invoke a service based on a particular binding.

The WSIFOperation interface is the run-time representation of an operation. This interface provides methods to create input, output, and fault messages, and to invoke the operation.

Here is the Javadoc for the WSIFOperation interface.

#### **createInputMessage, createOutputMessage and createFaultMessage**

These are factory methods to create the messages required by the invocation methods. All invocation methods require an input message.

#### **executeRequestResponseOperation**

This method invokes “In Out” operations.

#### **executeInputOnlyOperation**

This method invokes “In only” operations.

#### **executeRequestResponseOperation**

If this method is used for invocation, then an output and a fault message are instantiated and passed on the call to the method. If the method returns true, then the output message contains the response message. If the message returns false, then a fault occurred and is returned in the fault message.

#### **executeRequestResponseAsync**

This method allows “In Out” operations to be invoked with the reply handled using an alternate thread. Use of this method is discussed further in WSIFOperation - Asynchronous interactions.

#### **setContext and getContext**

Use of these methods is discussed in WSIFOperation - Context.

All of the **executeNnnn** methods fail with an exception if there is an error in processing the request in the WSIF provider.

Setting the timeouts for synchronous and asynchronous operations is discussed in WSIFOperation - Synchronous and asynchronous timeouts.

**WSIFOperation - Context:** Although WSDL does not define context, a number of uses of the Web Services Invocation Framework (WSIF) require the ability to pass context to the port that is invoking the service. For example, a SOAP over HTTP port might require an HTTP user name and password. This

information is specific to the invocation, but is not a parameter of the service. In general, context is defined as a set of name-value pairs. However, because Web services tend to define the types of data using XML schema types, WSIF represents the name-value pairs of the context using the same representation that WSIFMessage classes use; that is a set of named parts, each of which equates to an instance of an XML schema type.

You use the WSIFOperation interface setContext and getContext methods to pass context information to the binding. The port implementation can use this context, for example to update a SOAP header. There is no definition of how a port can utilize the context.

The parameter of the setContext and getContext methods is a WSIFMessage interface, and this interface has named parts defining the context information. The WSIFConstants class defines constants for the part names that can be set in a context WSIFMessage interface.

The following code fragment shows how to set the user name and password for HTTP basic authentication:

```
// set a basic authentication header
WSIFMessage headers = new WSIFDefaultMessage();
headers.setObjectPart(WSIFConstants.CONTEXT_HTTP_USER, "user name");
headers.setObjectPart(WSIFConstants.CONTEXT_HTTP_PSWD, "password");
operation.setContext(headers);
```

The WSIFOperation interface ignores context parts that it does not support. For example, the previous code is ignored by the WSIF Java provider.

The WSIFConstants class includes the following constants that can be used for context part names:

- CONTEXT\_HTTP\_USER
- CONTEXT\_HTTP\_PSWD
- CONTEXT\_SOAP\_HEADERS

The HTTP header values are expected to be of type String, and the SOAP header value is expected to be of type java.util.List, which should contain entries of type org.w3c.dom.Element.

***WSIFOperation - Asynchronous interactions reference:*** The Web Services Invocation Framework (WSIF) supports asynchronous operation. In this mode of operation, the client puts the request message as part of one transaction, and carries on with the thread of execution. The response message is then handled by a different thread, with a separate transaction. The SOAP over JMS and native JMS providers are the only WSIF providers that currently support asynchronous operation.

The WSIFPort class uses the supportsAsync method to test if asynchronous operation is supported.

An asynchronous operation is initiated with the WSIFOperation interface executeRequestResponseAsync method. This method lets a Remote Procedure Call (RPC) method be invoked asynchronously. The method returns before the operation is completed, and the thread of execution continues.

The response to the asynchronous request is processed by the WSIFOperation interface fireAsyncResponse or processAsyncResponse methods.

To initiate the request, there are two forms of the executeRequestResponseAsync method:

```
public WSIFCorrelationId executeRequestResponseAsync
 (WSIFMessage input, WSIFResponseHandler handler)
```

and

```
public WSIFCorrelationId executeRequestResponseAsync (WSIFMessage input)
```

**executeRequestResponseAsync(WSIFMessage input, WSIFResponseHandler handler)**

This method takes an input message and a WSIFResponseHandler handler. The handler is



invoked on another thread when the operation completes. When using this method the client listener calls the `fireAsyncResponse` method, which then calls the `WSIFResponseHandler` interface `executeAsyncResponse` method. Here is the Javadoc for the `WSIFResponseHandler` interface.

#### **executeRequestResponseAsync(WSIFMessage input)**

This method only takes an input message, and does not use a response handler. The client listener processes the response by calling the `WSIFOperation` interface `processAsyncResponse` method. This process updates the `WSIFMessage` output and fault messages with the result of the request.

WSIF supports correlation between the asynchronous request and response. When the request is sent, the `WSIFOperation` object is serialized into the `WSIFCorrelationService` object. The `executeRequestResponseAsync` methods return a `WSIFCorrelationId` object which identifies the serialized `WSIFOperation` object. The client listener can use this to match a response to a particular request.

The correlation service is located with the `WSIFCorrelationServiceLocator` class `getCorrelationService()` method in the `org.apache.wsif.utils` package.

In a managed container a default correlation service is defined in the default Java Naming and Directory Interface (JNDI) namespace using the name: `java:comp/wsif/WSIFCorrelationService`. If this correlation service is not available, then WSIF uses the `WSIFDefaultCorrelationService`.

Here is the Javadoc for the `WSIFCorrelationService` interface.

and this is the correlator ID:

```
public interface WSIFCorrelator extends Serializable {
 public String getCorrelationId();
}
```

The client must implement its own response message listener or message data base so that it can recognize the arrival of response messages. This client implementation manages the correlation of the response message to the request and call of one of the asynchronous response processing methods. As an example of the requirement for a client listener, the following code fragment shows what can be in the `onMessage` method of a Java Messaging Service (JMS) listener:

```
public void onMessage(Message msg) {
 WSIFCorrelationService cs = WSIFCorrelationServiceLocator.getCorrelationService();
 WSIFCorrelationId cid = new JmsCorrelationId(msg.getJMSCorrelationID());
 WSIFOperation op = cs.get(cid);
 op.fireAsyncResponse(msg);
}
```

#### ***WSIFOperation - Synchronous and asynchronous timeouts reference:***

When you use the Web Services Invocation Framework (WSIF) with the Java Messaging Service (JMS) you can set timeouts for synchronous and asynchronous operations.

Default values for these timeouts are defined in the WSIF properties file:

```
maximum number of milliseconds to wait for a response to a synchronous request.
Default value if not defined is to wait forever.
wsif.syncrequest.timeout=10000

maximum number of seconds to wait for a response to an async request.
if not defined on invalid defaults to no timeout
wsif.asyncrequest.timeout=60
```



If you use these default values, a synchronous request (such as a WSIFOperation interface executeRequestResponseOperation method call) times out after ten seconds, and an asynchronous request (such as a WSIFOperation interface executeRequestResponseAsync method call) times out after sixty seconds.

**Note:**

The code that processes both of these timeout values uses milliseconds as its unit of time. The WSIFProperties class getAsyncTimeout method multiplies the wsif.asyncrequest.timeout value by 1000, to convert the value from seconds to milliseconds.

You can override these default values for a given request by setting a JMS property on the operation request with the <jms:property> and <jms:propertyValue> WSDL elements. Set the name of the property to be the name of the timeout from the WSIF properties file.

The following example sets synchronous requests to time out after two minutes (120 seconds):

```
<jms:propertyValue name="wsif.syncrequest.timeout" type="xsd:string" value="120000"/>
```

and the following example disables asynchronous timeouts (a value of zero means wait forever):

```
<jms:propertyValue name="wsif.asyncrequest.timeout" type="xsd:string" value="0"/>
```

When an asynchronous timeout expires, no listener or message data base waiting for the response is notified. The asynchronous timeout is only used to tell the correlation service that the stored WSIFOperation can be deleted.

---

## Chapter 11. Enabling the IBM WebSphere UDDI Registry

Welcome to the IBM WebSphere UDDI Registry.

Use the table of contents (on the left and below) to view the various topics for a specific product or technology. Select the topic you are interested in to either open documentation locally or find information about how to locate documentation.

- Terminology
- Definitions
- Overview of UDDI Registries
- Migrating from the IBM WebSphere UDDI Registry on WebSphere Application Server 4.0
- Installing the UDDI Registry Component
- Applying Service to the UDDI Registry in a Network Deployment and single Application Server environment
- Use of a remote DB2 Database
- Reinstalling the UDDI Registry Component
- Removing the UDDI Registry application from a deployment manager cell
- Removing the UDDI Registry application from a single appserver
- Configuring the UDDI Registry
- Administering the UDDI Registry
- The UDDI user console
- Custom Taxonomy Support in the UDDI Registry
- The SOAP Application Programming Interface
- The application programming interface
- The EJB Interface
- UDDI4J
- Problem determination
- Messages
- Samples
- Installation Verification Program (IVP)
- Reporting Problems with the IBM WebSphere UDDI Registry
- Feedback

---

### UDDI Registry terminology

The directory location of the WebSphere Application Server is referred to as **<AppServer-install-dir>** and the directory location of the WebSphere Deployment manager as **<DeploymentManager-install-dir>**. The default locations are:

#### Windows

**<AppServer-install-dir>**

C:\Progra~1\WebSphere\AppServer\

**<DeploymentManager-install-dir>**

C:\Progra~1\WebSphere\DeploymentManager\

#### Linux/Solaris/HP Platforms

**<AppServer-install-dir>**

/opt/WebSphere/AppServer/

**<DeploymentManager-install-dir>**

/opt/WebSphere/DeploymentManager/

#### AIX Platform

**<AppServer-install-dir>**

/usr/WebSphere/AppServer/

<DeploymentManager-install-dir>  
/usr/WebSphere/DeploymentManager/

#### z/OS Platform

<AppServer-install-dir>  
/WebSphere390/V5R0M0/AppServer/  
<DeploymentManager-install-dir>  
/WebSphere390/V5R0M0/DeploymentManager/

## UDDI Registry definitions

### bindingTemplate

Technical information about a service entry point and construction specifications.

### businessEntity

Information about the party who publishes information about a family of services.

### businessService

Descriptive information about a particular service.

### publisherAssertion

Information about a relationship between two parties, asserted by one or both.

### tModel

Short for technical model.

A tModel is a data structure representing a reusable concept, such as a Web service type, a protocol used by Web services, or a category system.

tModel keys within a service description are a technical "fingerprint" that you can use to trace the compatibility origins of a given service. They provide a common point of reference that allows you to identify compatible services.

tModels are used to establish the existence of a variety of concepts and to point to their technical definitions. tModels that represent value sets such as category, identifier, and relationship systems are used to provide additional data to the UDDI core entities to facilitate discovery along a number of dimensions. This additional data is captured in keyedReferences that reside in category Bags, identifierBags, or publisherAssertions. The tModelKey attributes in these keyedReferences refer to the value set that relates to the concept or namespace being represented. The keyValues contain the actual values from that value set. In some cases keyNames are significant, such as for describing relationships and when using the general keywords value set. In all other cases, however, keyNames are used to provide a human readable version of what is in the keyValue.

---

## An overview of IBM UDDI Registries

The Universal Description, Discovery and Integration (UDDI) specification defines a way to publish and discover information about Web services. The term 'Web service' describes specific business functionality exposed by a company, usually through an Internet connection, to allow another company, or its subsidiaries, or software program to use the service.

### Universal Business Registries (IBM UBR)

The IBM Universal Business Registry is one of a group of Web-based registries that expose information about a business or other entity and its technical interfaces (or APIs). These registries are run by multiple Operator Sites, and can be used by anyone who wants to make information available about one or more businesses or entities, as well as anyone who wants to find that information. Although there are Universal Business Registries (sometimes referred to as 'public UDDI registries') hosted worldwide, including one hosted by IBM, enterprises may wish to host their own internal registries behind their firewall to better manage their internal implementation of Web services.

For more detailed information about UDDI in general visit <http://www.uddi.org>

## IBM WebSphere UDDI Registry

The IBM WebSphere UDDI Registry is a directory for Web services that is implemented using the UDDI specifications. In contrast with the IBM UBR, this component of WebSphere Network Deployment is a product offering for companies or industries to implement.

A critical component of IBM's dynamic e-business infrastructure, IBM WebSphere UDDI Registry solves the problem of discovery of technical components for an enterprise and its partners by:

- Providing control, flexibility and confidentiality so that an enterprise can protect its e-business investments
- Increasing efficiency by making it easier to identify technical assets
- Leveraging existing infrastructures

For example, the IBM WebSphere UDDI Registry could be used in the following way within a large enterprise:

A company has a legacy application that provides telephone numbers and Human Resources (HR) information of employees. This is turned into a Web service and published to the registry. A developer in the same company needs to write an application for a procurement function that also needs to provide HR information to the supplier. The application should allow the supplier to have access to the employee account codes once the employee provides his name or serial number. Before Web Services, the developer had three choices:

1. Would not have known about the similar application
2. Knew about it but could not reuse due to technical barriers
3. Knew about it and reused only after significant time and negotiation

With UDDI, the developer can search for the "web service" and reuse the existing technical component in their new application for the supplier in a matter of minutes. The developer saves time and gets the application up and running sooner than they would have otherwise -- increasing efficiency and saving the company time and money. The IBM WebSphere UDDI Registry is the first version 2 standard-compliant UDDI registry for private enterprise work. The IBM WebSphere UDDI Registry:

- Supports the public UDDI V2.0 standard
- Leverages the proven, reliable WebSphere Application Server technology
- Uses a relational database, such as DB2, for its persistent store.

---

## Migrating from a previous version of IBM WebSphere UDDI Registry

### Migrating from the IBM WebSphere UDDI Registry on WebSphere Application Server 4.0

If you have previously installed the IBM WebSphere UDDI Registry V1.1 (or later refreshes) on WebSphere Application Server V4.0, take the following steps in order to migrate to the UDDI Registry that is available as part of WebSphere Application Server for Network Deployment.

1. If you have made any changes to the configuration properties in the file `uddi.properties`, which is located in the `properties` subdirectory of your WebSphere Application Server install tree, make a copy of this file (or make a note of all the changes), so that you can reapply the changes to the file after you have installed WebSphere Application Server V5.0.
2. If you have made any other configuration changes, for example to the UDDI user console style sheets or to the SOAP interface properties, make a note of them, and reapply them after upgrading to the new UDDI version.
3. Uninstall the IBM WebSphere UDDI Registry using Add/Remove Programs on Windows platforms, or `rpm -e IBMWebSphere-UDDI` on Unix platforms. This removes the application, but preserves the UDDI Registry database.
4. If you wish to continue using DB2 as the persistence store for the UDDI Registry, and you have any data in the UDDI Registry that you wish to preserve, you should not run the DB2 setup wizard to

create the DB2 version of the UDDI Registry database, but will instead be able to continue using the database that you already have. However, if you do run the wizard, you are prompted to keep the database or overwrite it.

5. Now follow the instructions on installing (or upgrading) IBM WebSphere Application Server and "Installing the UDDI Registry component".
6. After completing the installation procedure, you should edit the `uddi.properties` file to reflect any changes that you require to the configuration properties. You should **not** replace this file with your previous copy from the IBM WebSphere UDDI Registry V1.1 or later refreshes. You can also reapply any other configuration changes as necessary.
7. If you are migrating from Version 1.1 of the IBM WebSphere UDDI registry, then there have been a few minor changes to the EJB interface which means that you may need to modify your EJB client applications. If you are migrating from later refreshes of the IBM WebSphere UDDI Registry (such as Version 1.1.1), then you should not need to make any changes to your EJB clients.

The changes for Version 1.1 relate to saving a new service through the EJB interface either using **saveBusiness** or **saveService**. You should not set the **serviceKey** before calling **saveBusiness** or **saveService**.

In the IBM WebSphere UDDI Registry V1.1 the term 'service type' was used to refer to a 'technical model' or 'tModel', for example, in various of the panels in the UDDI user console. This term has now been replaced by the term 'technical model'.

### 5.1+ Migrating from IBM WebSphere Network Deployment V5.0x - federated nodes - to Version 5.1

For a Network Deployment configuration with federated nodes, IBM WebSphere Application Server supplies facilities to migrate to Version 5.1 from Version 5.0x either as part of the installation process, or as a post installation procedure. For instructions on how to do this refer to "Migrating to Network Deployment" in the Information Center.

As part of the migration process, IBM WebSphere Application Server will automatically install the new UDDI version supplied in Version 5.1, but will migrate the `uddi.properties` (from your previous Version 5.0x system) for use in your new setup. If you are using DB2 as your persistence store, the same database used in Version 5.0x can be used in Version 5.1. Ensure your version of DB2 is supported by IBM WebSphere Application Server Version 5.1.

If you are using Cloudscape as your persistence store, you can also use the Version 5.0x database in Version 5.1 but you must review the Cloudscape instructions elsewhere in this Information Center for guidance on upgrading to a supported Cloudscape version.

By default, the newly migrated Application Server contains a JDBC provider that points to your 'old' pre-migration UDDI20 Cloudscape database (which is in your pre-migration WebSphere filing system). This is acceptable, however you may wish to copy your old UDDI20 database to the newly migrated Application Server filing system (in `<NewAppServer-install-dir>/bin/UDDI20`

In either case you will also need to make some updates to your JDBC provider to complete the migration, before starting the new Application Server. This can be done via the administration console.

- Update the JDBC provider resource as follows:
  - update the Classpath to point to the new location for `db2j.jar` which is `<NewAppServer-install-dir>/cloudscape/lib/db2j.jar`
- update the version 4 datasource for this
  - update the 'database name' path to point to the new location in the new Application Server, but only if you have moved it.
- update the custom properties
  - update the property 'database name' to point to the new location in the new Application Server, but only if you have moved it.

- add the new custom property 'connectionAttributes' with a value of 'upgrade=true'.

If you have made other configuration changes to the UDDI user console style sheets for example, or other SOAP interface properties, you will have to manually reapply them.

### **5.1 + Migrating from IBM WebSphere Application Server Version 5.0x to IBM WebSphere Application Server 5.1 in a single application server only environment**

In contrast with the Network Deployment migration, which is mostly automatic, some manual steps are required to upgrade from a WebSphere 5.0x single application server configuration to a WebSphere 5.1 single application server configuration.

In this case it is assumed that you have installed WebSphere Network Deployment to obtain the UDDI application and have installed UDDI into a standalone application server. It is assumed you wish to migrate your 5.0x application server only.

Refer to the migration instructions elsewhere in this Information Center (see "Migrating to Network Deployment") but here are the general steps you will need to follow:

- Stop the WebSphere 5.0x application server containing the UDDI
- Install IBM WebSphere Application Server Version 5.1 Network Deployment on to your system of choice. This will be needed to obtain the new UDDI application and associated files.
- Install IBM WebSphere Application Server Version 5.1 Application Server selecting the migration option. Once completed you will have a new 5.1 Application Server, with your configuration (other applications) migrated. This will include the UDDI application from your old 5.0x system.
- Copy across the new UDDI application and associated files from the **new** Version 5.1 Network Deployment to your new 5.1 Application Server. This is effectively following the same instructions as supplied in Installing and setting up a UDDI Registry, but in summary you will need to copy the following files and directories:
  - uddi.ear to **<NewAppServer-install-dir>\installableApps**
  - uddiresourcebundles.jar to **<NewAppServer-install-dir>\lib**
  - setupuddimessages.jar to **<NewAppServer-install-dir>\lib**
  - uddicloudscapeuserfunc.jar to **<NewAppServer-install-dir>\lib**
  - the UDDIReg/scripts directory to **<NewAppServer-install-dir>**

The uddi.properties file, including any modifications you have made are copied across by the migration process and requires no further action.

- **Database migration consideration**
  - **DB2 users.** If you are using DB2 as your persistence store, the same database used in Version 5.0x can be used in Version 5.1, but you must review the supporting software information to confirm your DB2 version is a supported level for IBM WebSphere Application Server Version 5.1
  - **Cloudscape users.** If you are using Cloudscape as your persistence store and want to continue to use the database (and the data in it) from your old WebSphere, copy your old UDDI20 to the newly migrated Application Server filing system (<NewAppServer-install-dir>/bin/UDDI20).
  - Start the new application server and run the appropriate appserversetupuddi.jacl command (for DB2 or Cloudscape) ensuring the paths refer to your new Version 5.1 Application Server. This installs the new Version 5.1 UDDI into your server. For details of the syntax of the appropriate commands see Setting up the UDDI Registry to use Cloudscape in a single Application Server or Setting up the UDDI Registry to use DB2 in a single Application Server and then return to this point.
- If you have made other configuration changes, to the UDDI user console style sheets for example or other SOAP interface properties, manually reapply them.
- Restart, and use, the new Application Server.



---

## Installing and setting up a UDDI Registry

If you wish to use the UDDI User Console using Internet Explorer as your Web browser, and using SSL, you must use Internet Explorer V5.5 with SP2 and security fix Q321232 (which must be applied in that order), or later.

If you are migrating from IBM WebSphere UDDI Registry running on an earlier version of WebSphere , read Migrating from a previous version of IBM WebSphere before installing the new product.

### Choice of database product to be used as the persistence store

The UDDI Registry application can use either DB2 or Cloudscape as the persistence store for the registry data.

- If you plan to use the UDDI Registry in production you must use DB2 as your persistence store.
- If you plan to use the UDDI Registry for development and testing purposes, you can also choose to use Cloudscape as your persistence store. Note that Cloudscape is not intended for production purposes.

### **Cloudscape Restriction**

Cloudscape Network Server Version 5.1 requires a WebSphere Version 5 datasource to utilize the multiple connection features. As IBM WebSphere UDDI Registry uses a WebSphere Version 4 datasource, this precludes other connections to the Cloudscape database when the UDDI Registry application is in the started state.

### Steps for this task

As part of the installation of the IBM WebSphere Application Server with Network Deployment option, you are given the option to install the "UDDI Registry", which is shown under Web Services. Having selected the UDDI Registry and installed the various files that make up the application, you have two choices as to the environment in which you run it:

1. Install the UDDI Registry application into the deployment manager cell using DB2 or Cloudscape as the database in which the UDDI Registry data will be held, selecting one of the application servers within the cell in which to run the UDDI Registry.
2. Install the UDDI Registry application directly into an application server using DB2 or Cloudscape as the database in which the UDDI Registry data will be held. If you choose this option, the application server on which you run the UDDI Registry must not be added into a deployment manager cell, as this would cause the file synchronization within the cell to remove the application.
3. It is also recommended that you execute `setupCmdline.bat` on Windows, or `./setupCmdline.sh` on Unix platforms, prior to executing any WebSphere commands.

In most cases you will probably choose option 1, and install the UDDI Registry into a deployment manager cell, but you might find that option 2, to install the UDDI Registry into a standalone application server, is useful for development or test purposes.

### **Note:**

1. Several WebSphere commands are used during the following procedures, some of which must execute on the DeploymentManager and some of which must execute on the target Application server. The instructions distinguish which is appropriate for each command. The WebSphere commands are in the bin subdirectory of the appropriate WebSphere install tree. To ensure correct operation of these commands, do one of the following:
  - Ensure that the appropriate bin subdirectory is in your path prior to executing the command
  - Change directory to the appropriate bin subdirectory
  - Fully qualify the path to the commands
2. It is recommended that you use the version of Java shipped with WebSphere and found in the `java/bin` subdirectory.



- It is also recommended that you execute `setupCmdline.bat` (on Windows) or `./setupCmdline.sh` (on Unix platforms) prior to executing any WebSphere commands.

The following table lists the UDDI Registry files, and the locations into which they are placed by the installation. The Location column shows the subdirectory under the WebSphere DeploymentManager install directory. For example, if you had installed IBM WebSphere Application Server with Network Deployment option onto a machine running Windows, and had used the default directory, then a location of `installableApps` would mean that the file had been placed into the `C:\Program Files\IBM\WebSphere\DeploymentManager\installableApps` directory. For Windows platforms, read the “/” directory separator in the location column as a “\” directory separation character.

Files	Purpose	Location
<code>uddi.ear</code>	The UDDI Registry application itself, which is packaged and runs as an enterprise application	<code>installableApps</code>
<code>uddi.properties</code>	Provides configuration properties for the UDDI Registry application	<code>properties</code>
<code>uddiresourcebundles.jar</code>	Contains system messages for the UDDI Registry application	<code>lib</code>
<code>uddicloudscapeuserfunc.jar</code>	Contains functions that are used by Cloudscape if the Cloudscape database is used with the UDDI Registry	<code>lib</code>
<code>setupuddi.jacl</code>	Administrative script to create a JDBC driver and datasource for the UDDI Registry, and to install the UDDI Registry application in a DeploymentManager Cell	<code>UDDIReg/scripts</code>
<code>setupuddimessages.jar</code>	Contains setup and install messages for the UDDI Registry application	<code>lib</code>
<code>removeuddi.jacl</code>	Administrative script to undo the effects of <code>setupuddi.jacl</code>	<code>UDDIReg/scripts</code>
<code>appserverremoveuddi.jacl</code>	Administrative script to undo the effects of <code>appserversetupuddi.jacl</code>	<code>UDDIReg/scripts</code>
<code>appserversetupuddi.jacl</code>	Administrative script to create a JDBC driver and datasource for the UDDI Registry, and to install the UDDI Registry application in a single, stand-alone, application server	<code>UDDIReg/scripts</code>
<code>SetupDB2UDDI.jar</code>	The 'UDDI DB2 Setup Wizard', to create and pre-load the UDDI Registry database if DB2 is to be used as the persistence store	<code>UDDIReg/scripts</code>
<code>UDDI20</code> (directory)	Cloudscape directory containing the UDDI Registry tables and pre-loaded data	<code>bin</code>
<code>uddiejbclient.jar</code>	Class library for use when writing an EJB client to access the UDDI Registry	<code>UDDIReg/ejb</code>
Various javadoc files	JAVADOC to describe the EJB interface to the UDDI Registry	<code>UDDIReg/ejb/javadoc</code>
<code>UDDITaxonomyTools.jar</code>	Provides tools for supporting custom taxonomies with the UDDI Registry	<code>UDDIReg/scripts</code>
<code>CustomTaxonomy.properties</code>	Provides configuration properties to be used the <code>UDDITaxonomyTools</code>	<code>UDDIReg/scripts</code>
<code>UDDIUtilityTools.jar</code>	Provides support for import/export of UDDI entities	<code>UDDIReg/scripts</code>
<code>UDDIUtilityTools.properties</code>	Provides configuration properties for the UDDI Utility Tools	<code>UDDIReg/scripts</code>


If you intend to run in a Deployment Manager Cell then complete the following task - **Installing the UDDI Registry into a deployment manager cell**

If you intend to run in a single WebSphere Application server, then complete the following task - **Installing the UDDI Registry into a single WebSphere Application Server**

Continue with Configuring the UDDI Registry.

## Installing the UDDI Registry into a deployment manager cell

The diagram following shows the configuration used for the Deployment Manager example configurations



that follow:

In this configuration, several nodes are federated to the Deployment Manager (dmgr) on myhome.headoffice.xyz.com, from which UDDI is deployed to the host myriad.headoffice.xyz.com.

These instructions assume that the installation has been performed into a clean environment. If you are installing into an existing deployment manager cell skip to step 6.

1. Install the WebSphere Application Server for Network Deployment package, and select Deployment Manager and UDDI Registry options (under Web Services).
2. Install one or more base application servers which will form the cell of servers. One of these should be the application server in which you plan to run an instance of the UDDI Registry. You can run more than one instance of the UDDI Registry within a cell of servers: the UDDI Registry application name is suffixed with the target node and server names to make it unique within the cell (See also "Advanced use of setupuddi.jacl"), but you can only run one UDDI instance within each application server.
3. Ensure that the target application server is stopped.
4. Run startManager (startManager.sh on Unix and Linux platforms) on the deployment manager node to start the deployment manager.
5. Run addNode (addNode.sh on Unix and Linux platforms) on each of the baseapplication server(s) to add it as a node into the cell. (How to run addNode is described elsewhere in the Information Center - see addNode command). For example: addnode myhome- where myhome is the IP name of your deployment manager host.
6. Copy the *uddiejbclient.jar* file and the EJB javadoc directory tree from the UDDIReg/ejb subdirectory of the deployment manager install tree onto any machine(s) where you will be creating EJB clients to access the UDDI Registry.
7. If you have any global configuration properties that are common to any UDDI Registries that you install into this cell, you can edit the *uddi.properties* file in the properties subdirectory of the

deployment manager install tree to set them up. (See the section on Configuring global UDDI properties for more details about the global configuration properties).

8. If required, edit the security permissions for the UDDI Registry application. You should only do so if you have a thorough understanding of Java 2 security issues, and the way in which security permissions are used by WebSphere. The permissions for the UDDI Registry application are set within the *was.policy* file, which is part of the *uddi.ear*. To see and change the contents of this file you should:

- a. On the deployment manager, copy the **uddi.ear** file from the installableApps subdirectory of the deployment manager install tree into a temporary directory.
- b. Un-jar the *uddi.ear* file (that is unpack *uddi.ear* using the 'jar -x' command).

For example:

```
jar -x uddi.ear
```

(This uses the jar command in the <DeploymentManager-install-dir>\java\bin subdirectory of the deployment manager, so you might need to fully qualify the path to the jar command.)

- c. You will find the *was.policy* file under the META-INF subdirectory that is created.

This will allow you to see the permissions which have been granted to the UDDI Registry application, and to make any changes that are necessary. Please note that if you make any errors in changing this file, then the UDDI Registry application might either fail to start, or will encounter errors when trying to execute UDDI requests.

- d. Re-jar the *uddi.ear* file using the jar command.

For example:

```
jar -cf uddi.ear .
```

**Note:** NOTE the space and the dot after *uddi.ear*)

(This uses the jar command in the <DeploymentManager-install-dir>\java\bin subdirectory of the deployment manager, so you might need to fully qualify the path to the jar command.)

- e. Copy the new *uddi.ear* back to the installableApps directory.
9. Note that if the target application server is running, this step will stop and restart it. If you are planning to use Cloudscape for the database in which the UDDI Registry will be held, please read the section "Setting up the UDDI Registry to use Cloudscape within a deployment manager cell" and then return to this point. If however, you plan to use DB2, then please refer to the section "Setting up the UDDI Registry to use DB2 within a deployment manager cell" and then return to this point.
10. Ensure that the UDDI Registry is configured appropriately for your installation, as described in the section on Configuring the UDDI Registry.
11. Start, or stop and restart, the target application server. This should also start the UDDI Registry application. If not, use the administrative console on the deployment manager to do so.  
On Unix and Linux platforms run the *db2profile* script before issuing the *startServer.sh* server1 command. This script is located within the DB2 instance's home directory under *SQLLIB* and you can invoke it by typing:

```
". /home/db2inst1/sqllib/db2profile"
```

**Note:** In the above example, notice that the '.' is followed by a single space character.

**Note:** On Unix and Linux platforms the DB2 user **must** have a *db2profile* at *\$HOME/sql/lib/db2profile*.

12. Your UDDI application is now ready to use. Go to the User Console section within this Information Center or any of the API sections.

### **Advanced use of setupuddi.jacl**

A number of symbols are defined at the top of the setupuddi.jacl script. These allow you to control the amount of logging that is performed, and to install multiple instances of the UDDI Registry within the same cell.

The symbols that you can edit are as follows:

- **logEnabled** - default setting is 1, which causes the progress of the script to be logged. Setting this symbol to 0 causes information logging to be suppressed, with only error messages being output.
- **overwriteExisting** - default setting is 1 which causes any existing installation of the UDDI Registry application to be overwritten. Setting this symbol to 0 would cause the existing installation to be left as is, but would allow other files used by the UDDI Registry to be updated. You are recommended to only change this setting under the guidance of IBM Service.
- **appName** - default setting is UDDIRegistry, which is the first part of the name used for the UDDI Registry application installed into the target server. To ensure uniqueness of application names within the cell, the full application name that will be used is `<appName>.<nodeName>.<server>`, where `<nodeName>` is the name of the target node and `<server>` is the name of the target server. You can change the first part of this (the `<appName>`) portion by changing the setting of this symbol before running setupuddi.jacl, although it is generally recommended that you do not change this value.

Continue with Configuring the UDDI Registry.

## Setting up the UDDI Registry to use Cloudscape within a deployment manager cell

If you plan to use Cloudscape for the database in which the UDDI Registry data will be held, perform this task to setup and install the UDDI Registry database to use the supplied Cloudscape database.

See "Choice of database product to be used as the persistence store" to decide which database product you should use as your persistence store before proceeding further with this task.

This task is part of a parent task: Installing the UDDI Registry into a deployment manager cell. You should complete this task at the appropriate step in the parent task.

This task configures Cloudscape on the host where you want to run the UDDI Registry. Cloudscape is supplied with WebSphere Application Server.

In this task you will invoke a script called setupuddi.jacl, specifying the target node and application server into which the UDDI Registry is to be deployed. If the target application server is running when you invoke setupuddi.jacl, the script stops the server and restarts the server after it has completed its operations.

1. Copy the UDDI20 directory tree from the bin subdirectory of the deployment manager tree into the bin subdirectory of the target application server's installation tree.
2. Create a JDBC driver and datasource to provide access to the UDDI20 Cloudscape database, and install the UDDI Registry application. This is done using the wsadmin tool, using as input the setupuddi.jacl script from the UDDIReg/scripts subdirectory of the Deployment Manager. Note that this script must be run on the deployment manager node.

You should either run this script from the UDDIReg/scripts subdirectory where it is located, or copy it to some other suitable directory. Note that the wsadmin command is located in the bin subdirectory of the deployment manager node. The syntax for calling this script for Cloudscape is:

```
wsadmin -f setupuddi.jacl
 deploymgrpath
 servername
 nodename
 discoveryURLprefix
 pathtdb
 > setupuddi.log
```

where

- *deploymgrpath* is the fully qualified pathname of the deployment manager install directory, specified using forward slashes regardless of platform; for example for Windows, this might be C:/Progra~1/WebSphere/DeploymentManager for Windows, or, for Unix platforms it might be /opt/WebSphere/DeploymentManager.
- *servername* is the name of the target server on which you wish to deploy the UDDI Registry, such as server1. Note the the server name entered is case sensitive.
- *nodename* is the name of the WebSphere node on which the target server runs. Note the the node name entered is case sensitive.
- *discoveryURLprefix* is the URL prefix to be used for discovery URLs. Typically this will be of the form http://<ip-address>:9080/uddisoap/ - an example of a discoveryURLprefix value might be http://mynode.mylocation.mycompany.com:9080/uddisoap/
- *pathtodb* is the path to the UDDI20 database within the bin subdirectory of your WebSphere AppServer installation, specified using forward slashes regardless of platform; for example for Windows, this might be C:/Progra~1/WebSphere/AppServer/bin/UDDI20 and for Unix platforms: /opt/WebSphere/AppServer/bin/UDDI20
- > *setupuddi.log* is an optional parameter to direct the output to a log file as opposed to the default (which is to the screen)

For example on Windows (shown here on multiple lines for publication):

```
wsadmin -f setupuddi.jacl "C:/Progra~1/WebSphere/DeploymentManager/" server1 myriad
"http://myriad.headoffice.xyz.com:9080/uddisoap/" "C:/Progra~1/WebSphere/Appserver/bin/UDDI20"
```

or, on Unix platforms (shown here on multiple lines for publication):

```
./wsadmin.sh -f setupuddi.jacl "/opt/WebSphere/DeploymentManager/" server1 myriad
"http://myriad.headoffice.xyz.com:9080/uddisoap/" "/opt/WebSphere/Appserver/bin/UDDI20"
```

installs the UDDI Registry application into the server server1 running on node myriad, and sets it up to access the Cloudscape UDDI20 database located in the bin subdirectory of the application server.

The setupuddi.jacl script:

- Creates a JDBC driver named *UDDI.JDBC.Driver.<nodeName>.<server>* and a datasource named *UDDI.Datasource.<nodeName>.<server>* (where *<nodeName>* is the name of the target node and *<server>* is the name of the target server, and will replace any existing driver and datasource of that name.
- Checks whether the UDDI Registry application is already installed and, if so, stop it and uninstall it.
- Updates the uddi.properties configuration property file to configure the discoveryURLprefix value that you have specified and set the persister property as 'Cloudscape', and place this file into the location *config/cells/<currentcell>/nodes/<nodename>/servers/<servername>/uddi.properties*.
- Places a number of files that are needed by the UDDI Registry into the WebSphere configuration repository, and updates the ws.ext.dirs list to reference these files.
- Installs the UDDI Registry.

This script deploys the UDDI Registry into the configuration under the deployment manager, and then do a Synchronization to install it into the specified server.

**Note:** The setup script, setupuddi.jacl, cannot be used to install the UDDI Registry application into a clustered application server. It is possible to cluster the UDDI Registry application by installing UDDI into an unclustered application server using the setup script, and then cluster that application server.

Return to the next step in the parent task Installing the UDDI Registry into a deployment manager cell.

## Setting up the UDDI Registry to use DB2 within a deployment manager cell

To decide which database product you should use as your persistence store, see "Choice of database product to be used as the persistence store".

This task is part of a parent task: Installing the UDDI Registry into a deployment manager cell. You should complete this task at the appropriate step in the parent task.

If you plan to use DB2 for the database in which the UDDI Registry will be held, ensure that the correct prerequisite fix packs have been applied as listed at <http://www-3.ibm.com/software/webservers/appserv/doc/latest/prereq.html> otherwise the startup of the UDDI DB2 setup wizard will fail.

If you plan to use DB2 for the database in which the UDDI Registry data will be held, use this task to create and load the UDDI Registry database using DB2, and to setup and install the UDDI Registry application to use the DB2 database.

This task uses the UDDI DB2 Setup Wizard to configure DB2 on the system where you want to run the UDDI Registry. Before starting this task, ensure that DB2 is installed and running on that system.

Copy the UDDIReg directory tree from the deployment manager to the target application server where DB2 will run.

The following steps should be carried out on the system on which the target application server is located (referred to as the 'target system').

In this task you will invoke a script called `setupuddi.jacl`, specifying the target node and application server into which the UDDI Registry is to be deployed. If the target application server is running when you invoke `setupuddi.jacl`, the script stops the server and restarts the server after it has completed its operations.

Before starting this task, ensure that you have created an appropriate DB2 userid and password. This same userid and password must be used throughout the following steps where the DB2 userid and password is requested.

1. On Windows, and if using DB2 Version 7, ensure that since installing DB2 you have run the `usejdbc2.bat` command file. For more information about this, see "Application Building Guide" in the DB2 documentation.
2. Create and load the UDDI Registry database, called UDDI20.

**Note:** If you are migrating from an earlier version of the UDDI Registry, and your UDDI20 DB2 database already exists, skip this step unless you want to replace the existing database with a new UDDI20 DB2 database. If you do choose to replace an existing database, all of your existing UDDI data will be lost. **Important:** if you do choose to replace the existing UDDI20 database, there must not be any applications or users accessing the database at the time that you run the UDDI DB2 setup wizard.

**Note:** It is recommended you use the `SetupDB2UDDI.jar` command from IBM WebSphere Application Server Version 5.0.1 or later. This is essential for non-English users.

**Note:** If you have a copy of the file `SetupDB2UDDI.jar` in your application server directory, then the application of the base and Network Deployment PTFs will not update `SetupDB2UDDI.jar` in your application server directory. You must apply the PTF for Network Deployment to your DeploymentManager file structure to update the `SetupDB2UDDI.jar` located there (in the `/UDDIReg/scripts` subdirectory), and then manually copy this jar file to any application server you may wish to run it on.

If you are planning to use a remote DB2 system on another host machine, copy the `SetupDB2UDDI.jar` file to the remote system and run it on that system to create and load the UDDI Registry database following the instructions within this step and continue with the next step (which states "If using a remote DB2 system on another host machine..") on the local host.

The UDDI DB2 Setup Wizard used in this task will prompt you to provide the DB2 userid and password under which the UDDI Registry database will be created and subsequently accessed.

Information on how to do this and where to obtain the PTF can be found in "Installing interim fixes, cumulative fixes, and fix packs" in the Information Center.



On Windows this should be a userid and password with administrative privileges. On Unix platforms, you should supply the userid and password of the DB2 instance in which you wish the database to be created.

To create the database you use the UDDI DB2 setup wizard, which is supplied as a jar file called SetupDB2UDDI.jar in the UDDIReg/scripts subdirectory, by following these steps:

- a. Change directory to the directory containing the file SetupDB2UDDI.jar (that is, either the UDDIReg/scripts directory in which it is supplied, or a directory on the target system into which you have copied it).
- b. To run the wizard, first ensure that you have access from your command line to the JVM supplied with WebSphere:

- On Windows, in a command window, a user with administrative privileges should type the following command:

```
<AppServer-install-dir>\bin\setupcmdline.bat
```

For example,

```
C:\Program Files\WebSphere\AppServer\bin\setupcmdline.bat
```

- On Unix platforms, at a command line type one of the following commands:

- If you are using bash, then as the root user run

```
. <AppServer-install-dir>/bin/setupCmdLine.sh
```

- If you are using csh, then as the root user run

```
source . <AppServer-install-dir>/bin/setupCmdLine.sh
```

- c. Temporarily set your path by typing:

- On Windows:

```
set path=%WAS_PATH%;%path%
```

- On Unix or Linux platforms:

```
export PATH=$WAS_PATH:$PATH
```

- d. In the same command window, start the UDDI DB2 setup wizard by entering one of the following commands:

- To start a graphical user interface, type

```
java -jar SetupDB2UDDI.jar
```

- To start a text mode interface, type:

```
java -jar SetupDB2UDDI.jar -console
```

**Note:** The above commands will fail unless your DB2 installation has the correct fix packs applied as listed on the prerequisite page found at <http://www-3.ibm.com/software/webservers/appserv/doc/latest/prereq.html>

- e. Follow the prompts to work through the wizard panels or command prompts.

- f. If necessary, check the log files for the wizard. A log file called UDDIloadDB.log is written to the directory from which the wizard is run (but note that on Windows platforms, if you have decided not to overwrite an existing UDDI20 database, then this fact is not logged, and the log file is not be created).

3. If using a remote DB2 system on another host machine, refer to "Use of a remote DB2 database" and then return to this point and continue with the following instructions.
4. Create a JDBC driver and datasource to provide access to the UDDI20 DB2 database, and install the UDDI Registry application. This is done using the wsadmin tool, using as input the setupuddi.jacl script from the UDDIReg/scripts subdirectory of the deployment manager. This script must be run on the deployment manager node.

Either run this script from the UDDIReg/scripts subdirectory where it is located, or copy it to some other suitable directory. Note that the wsadmin command is located in the bin subdirectory of the deployment manager node. The syntax for this script for DB2 is:

```
wsadmin -f setupuddi.jacl
 deploymgrpath
 servername
```



```

nodename
discoveryURLprefix
dbname
db2userid
db2password
db2ziplocation
> setupuddi.log

```

where:

- *deploymgrpath* is the fully qualified pathname of the deployment manager install directory, specified using forward slashes regardless of platform; for example for Windows, this might be `c:/Progra~1/WebSphere/DeploymentManager` and, for Unix platforms it might be `/opt/WebSphere/DeploymentManager/`
- *servername* is the name of the target application server on which you wish to deploy the UDDI Registry, such as `server1`. The server name is case sensitive.
- *nodename* is the name of the WebSphere node on which the target application server runs. Typically, this is the same as the machine name. The node name is case sensitive.
- *discoveryURLprefix* is the URL prefix to be used for discovery URLs. Typically this will be of the form `http://<ip-address>:9080/uddisoap/` so an example of a *discoveryURLprefix* value might be `http://mynode.mylocation.mycompany.com:9080/uddisoap/`
- *dbname* is the name of the UDDI Registry database under DB2. For this parameter, you should specify `UDDI20`.

**Note:** If a remote DB2 system is being used the *dbname* stated here must be the alias created when access to the remote database was set up on the local system

- *db2userid* and *db2password* are a valid DB2 userid and password with administrative privileges
- *db2ziplocation* is the path to the `db2java` zip file on your system, specified using forward slashes regardless of platform; for example for Windows, this might be `C:/Progra~1/SQLLIB/java/db2java.zip` or, for Unix platforms it might be `/home/db2inst1/sqllib/java12/db2java.zip`
- `> setupuddi.log` is an optional parameter to direct the output to a log file as opposed to the default (which is to the screen)

For example:

On Windows the command is (shown here on multiple lines for publication):

```

wsadmin -f setupuddi.jacl "C:/Progra~1/WebSphere/DeploymentManager/" server1 myriad
"http://myriad.headoffice.xyz.com:9080/uddisoap/" UDDI20 db2admin secretpwd
"C:/Progra~1/SQLLIB/java/db2java.zip" > setupuddi.log

```

On Unix platforms the command is (shown here on multiple lines for publication):

```

wsadmin.sh -f setupuddi.jacl "/opt/WebSphere/DeploymentManager/" server1 myriad
"http://myriad.headoffice.xyz.com:9080/uddisoap/" UDDI20 db2admin secretpwd
"/home/db2inst1/sqllib/java12/db2java.zip" > setupuddi.log

```

This installs the UDDI Registry application into the server `server1` running on node `myriad`, and set it up to access the DB2 `UDDI20` database using the userid `'db2admin'` and password `'secretpwd'`.

The `setupuddi.jacl` script:

- Creates a JDBC driver named `UDDI.JDBC.Driver.<nodeName>.<server>` and a datasource named `UDDI.Datasource.<nodeName>.<server>` (where `<nodeName>` is the name of the target node and `<server>` is the name of the target server), and replaces any existing driver and datasource of that name.
- Checks whether the UDDI Registry application is already installed and, if so, stop it and uninstall it.
- Updates the `uddi.properties` configuration file to configure the `discoveryURLprefix` value that you have specified, and to set the `persist` property as `'DB2'`, and places this file into the location `config/cells/<currentcell>/nodes/<nodename>/servers/<servername>/uddi.properties`.
- Places a number of files that are needed by the UDDI Registry into the WebSphere configuration repository, and update the `ws.ext.dirs` list to reference these files.
- Installs the UDDI Registry.

**Note:** The setup script, `setupuddi.jacl`, cannot be used to install the UDDI Registry application into a clustered application server. It is possible to cluster the UDDI Registry application by installing UDDI into an unclustered application server using the setup script, and then cluster that application server.

Return to the next step in the parent task Installing the UDDI Registry into a deployment manager cell.

## Installing the UDDI Registry into a single appserver

If you intend to run in a single WebSphere Application server, then complete the following task.

When you select the UDDI Registry option, then the installation will place all files that are needed to run a UDDI Registry onto the deployment manager install tree on the machine on which you are installing IBM WebSphere Application Server for Network Deployment.

To be able to run the UDDI Registry in a single application server instance in your network space you must copy these files over to the application server and then deploy the UDDI Registry. You can do this as follows:

1. Stop the application server on which you plan to run the UDDI Registry; for example, using the command `stopServer server_name` (`stopServer.sh` on Unix and Linux Platforms)
2. Copy the `uddi.ear` file from the `installableApps` subdirectory of the deployment manager install tree into the `installableApps` subdirectory of the target application server's install tree.
3. Copy the `uddi.properties` file from the `properties` subdirectory of the deployment manager install tree into the `properties` subdirectory of the target application server's install tree. In a subsequent step, you configure the UDDI Registry using the properties in the `uddi.properties` file.
4. Copy both the `uddiresourcebundles.jar` and the `setupuddimessages.jar` files from the `lib` subdirectory of the deployment manager install tree into the `lib` subdirectory of the target application server's install tree.
5. Optionally, if you are going to write or run code that uses the EJB interface to UDDI on another machine, then copy the `uddiejbclient.jar` file and the EJB javadoc directory tree from the `UDDIReg/ejb` subdirectory of the deployment manager install tree onto a location of your choice on any machines where you will be creating EJB clients to access the UDDI Registry.
6. Configure database support for the UDDI Registry database, in which the UDDI Registry will be held. To do this, complete one of the following tasks then return this point:
  - Setting up the UDDI Registry to use Cloudscape in a single AppServer
  - Setting up the UDDI Registry to use DB2 in a single AppServer

**Note:** If you set up the UDDI Registry application with a JDBC driver and datasource that reference Cloudscape, but set the `persist` property in `uddi.properties` to specify DB2, **or vice versa**, then some unexpected behavior will result, such as a fatal error on deleting an entity. If this happens, you should check that the above details are not in conflict. This only applies to a UDDI Registry installation on a single appserver.

7. Ensure that the UDDI Registry is configured appropriately for your installation, as described in the section on Configuring the UDDI Registry.
8. Stop and then restart the application server. On Unix and Linux platforms you must remember to run the `db2profile` script before issuing the `startServer.sh` command. This script is located within the DB2 instance's home directory under `SQLLIB` and can be invoked, for example, by typing:

```
. /home/db2inst1/sqllib/db2profile
```

**Note:** In the above example, it should be noticed that the `'.'` is followed by a single space character.

**Note:** On Unix and Linux platforms the DB2 user **must** have a `db2profile` at `$HOME/sqllib/db2profile`.

Continue with Configuring the UDDI Registry.

## Setting up the UDDI Registry to use Cloudscape in a single application server

To decide which database product you should use as your persistence store, see "Choice of database product to be used as the persistence store".

If you plan to use Cloudscape for the database in which the UDDI Registry data is held, use this task to set up and install the UDDI Registry application to use the supplied Cloudscape database.

This task configures Cloudscape on the host where you want to run the UDDI Registry.

This task, to configure Cloudscape for the UDDI Registry database, is part of a parent task Installing and Setting up a UDDI Registry. You should complete this task at the appropriate step in the parent task.

1. Copy the UDDI20 directory tree from the bin subdirectory of the deployment manager tree into the bin subdirectory of the target application server's install tree.
2. Copy the `uddicloudscapeuserfunc.jar` file from the lib subdirectory of the deployment manager install tree to the lib subdirectory of the target application server's install tree.
3. Ensure that the `persist` property in the `uddi.properties` file is set to `persist=Cloudscape`
4. Copy the `appserversetupuddi.jacl` script from the `UDDIReg/scripts` subdirectory of the deployment manager install tree to the WebSphere Application Server bin subdirectory.
5. Change directory to the WebSphere Application Server bin subdirectory.
6. Start the application server on which the UDDI Registry is to run.

For example, enter:

```
startServer server1 for Windows
. /startServer.sh server1 for Unix platforms
```

7. Create a JDBC driver and datasource to provide access to the UDDI20 Cloudscape database, and install the UDDI Registry application. To do this run the `wsadmin` tool with the script `appserversetupuddi.jacl` as input, on the target application server, using the following command syntax: (You should either run this script from the `UDDIReg/scripts` subdirectory where it is located, or copy it to some other suitable directory. Note that the `wsadmin` tool is located in the WebSphere bin subdirectory.)

```
wsadmin -f appserversetupuddi.jacl
 uddi-ear-location
 servername
 nodename
 WebSphere-lib-subdirectory
 cloudscapedbname
 > setupuddi.log
```

Where:

- `uddi-ear-location` is the fully-qualified path to the `uddi.ear` file in the `installableApps` subdirectory, specified **using forward slashes regardless of platform**. For example, on Windows:  
C:/Progra~1/WebSphere/AppServer/installableApps/uddi.ear  
and on Unix platforms:  
/opt/WebSphere/AppServer/installableApps/uddi.ear
- `servername` is the name of the application server on which the UDDI registry is to run; for example: `server1`. Note that the server name entered is case sensitive.
- `nodename` is the name of the WebSphere node on which the application server, `servername`, is running. Typically this is the machine name. Note that the node name entered is case sensitive.
- `WebSphere-lib-subdirectory` is the fully-qualified path to the WebSphere Application Server lib subdirectory, specified using forward slashes regardless of platform. For example on Windows:  
C:/Progra~1/WebSphere/AppServer/lib  
and on Unix platforms:

- `/opt/WebSphere/AppServer/lib`
- *cloudscapedbname* is the fully-qualified path to the UDDI20 database within the bin subdirectory of your WebSphere AppServer installation, specified using forward slashes regardless of platform. For example on Windows:
  - `C:/Program~1/WebSphere/AppServer/bin/UDDI20`
 and on Unix platforms
  - `/opt/WebSphere/AppServer/bin/UDDI20`
- `> setupuddi.log` is an optional parameter to direct the output to a log file as opposed to the default (which is to the screen)

The `appserversetupuddi.jacl` script completes the following actions:

- a. Creates a JDBC driver named `UDDI.JDBC.Driver.<nodeName>.<server>` and a datasource named `UDDI.Datasource.<nodeName>.<server>` (where `<nodeName>` is the name of the target node and `<server>` is the name of the target server, and will replace any existing driver and datasource of that name.
- b. Checks whether the WebSphere UDDI Registry application is already installed and, if so, stop the application and uninstall it.
- c. Installs the WebSphere UDDI Registry, then starts it.

**Note:** The setup script, `appserversetupuddi.jacl`, cannot be used to install the UDDI Registry application into a clustered application server. It is possible to cluster the UDDI Registry application by installing UDDI into an unclustered application server using the setup script, and then clustering that application server.

Return to the next step in the parent task Installing the UDDI Registry into a single appserver.

## Setting up the UDDI Registry to use DB2 in a single application server

To decide which database product you should use as your persistence store, see "Choice of database product to be used as the persistence store".

This task is part of a parent task Installing and setting up a UDDI Registry. You should complete this task at the appropriate step in the parent task.

If you plan to use DB2 for the database in which the UDDI Registry is held, ensure that the correct prerequisite fix packs have been applied as listed at <http://www-3.ibm.com/software/webservers/appserv/doc/latest/prereq.html> otherwise the startup of the UDDI DB2 setup wizard will fail.

If you plan to use DB2 for the database in which the UDDI Registry data will be held, use this task to create and load the UDDI Registry database using DB2, and to setup and install the UDDI Registry application to use the database.

This task uses the UDDI DB2 setup wizard to configure DB2 on the system where you want to run the UDDI Registry. Before starting this task, ensure that DB2 is installed and running on that system.

Copy the UDDIReg directory tree from the deployment manager to the target application server where DB2 will run.

The following steps should be carried out on the system on which the target application server is located (referred to below as the 'target system').

1. On Windows, if you are using DB2 Version 7, ensure that since installing DB2 you have run the `usejdbc2.bat` command file. For more information about this, see "Application Building Guide" in the DB2 documentation.
2. Create and load the UDDI Registry database, called UDDI20.

**Note:** If you are migrating from an earlier version of the UDDI Registry, and your UDDI20 DB2 database already exists, you should skip this step unless you want to replace the existing database with a new UDDI20 DB2 database. If you do choose to replace an existing database all of your existing UDDI data will be lost. **Important:** If you do choose to replace the existing UDDI20 database, then there must not be any applications or users accessing the database at the time that you run the UDDI DB2 setup wizard.

**Note:** It is recommended you use the `setupDB2UDDI.jacl` command from IBM WebSphere Application Server Version 5.0.1 or later. This is essential for non-English users.

**Note:** If you have a copy of the file `SetupDB2UDDI.jar` in your application server directory, then the application of the base and Network Deployment PTFs will not update `SetupDB2UDDI.jar` in your application server directory. You must apply the PTF for Network Deployment to your `DeploymentManager` file structure to update the `SetupDB2UDDI.jar` located there (in the `/UDDIReg/scripts` subdirectory), and then manually copy this jar file to any application server you may wish to run it on.

If you are planning to use a remote DB2 system on another host machine, copy the `SetupDB2UDDI.jar` file to the remote system and run it on that system to create and load the UDDI Registry database following the instructions within this step and continue with the next step (which states "If using a remote DB2 system on another host machine ....") on the local host.

The UDDI DB2 setup wizard used in this task prompts you to provide the DB2 userid and password under which the UDDI Registry database is created and subsequently accessed. Before starting this task, ensure that you have created an appropriate DB2 userid and password. This same userid and password must be used throughout the following steps where the DB2 userid and password is requested.

On Windows this should be a userid and password with administrative privileges. On Unix platforms, you should supply the userid and password of the DB2 instance in which you wish the database to be created.

To create the database you use the UDDI DB2 setup wizard:

- a. Change directory to the directory containing the file `SetupDB2UDDI.jar` (that is, either the `UDDIReg/scripts` directory in which it is supplied, or a directory on the target system into which you have copied it).
- b. To run the wizard, you need to first ensure that you have access from your command line to the JVM supplied with WebSphere. This is done as follows:
  - On Windows, in a command window, a user with administrative privileges should type the following command:

```
<AppServer-install-dir>\bin\setupcmdline.bat
```

For example, `C:\Progra~1\WebSphere\AppServer\bin\setupcmdline.bat`
  - On Unix platforms, enter one of the following commands as the root user:
    - If you are using bash:

```
. <AppServer-install-dir>/bin/setupCmdLine.sh
```
    - If you are using csh:

```
source /<AppServer-install-dir>/bin/setupCmdLine.sh
```
- c. Temporarily set your path by typing:
  - On Windows:

```
set path=%WAS_PATH%;%path%
```
  - On Unix or Linux platforms:

```
export PATH=$WAS_PATH:$PATH
```
- d. In the same command window, start the UDDI DB2 setup wizard by issuing one of the following commands:
  - To start a graphical user interface, enter:

```
java -jar SetupDB2UDDI.jar
```

- To start a text mode interface, enter:

```
java -jar SetupDB2UDDI.jar -console
```

**Note:** The above commands will fail unless your DB2 installation has the correct fix packs applied as listed on the prerequisite page found at <http://www-3.ibm.com/software/webservers/appserv/doc/latest/prereq.html>

- e. Follow the prompts to work through the wizard panels or command prompts.
- f. If necessary, check the log files for the wizard. A log file called UDDIloadDB.log is written into the directory from which the wizard is run (but note that, on Windows platforms, if you have decided not to overwrite an existing UDDI20 database, this fact is not logged, and the log file is not created).
3. If using a remote DB2 system on another host machine, refer to "Use of a remote DB2 database" and then return to this point and continue with the next step.
4. Ensure that the persist property in the uddi.properties file is set to persist=DB2.
5. On Unix, run the db2profile script to set up the environment for the DB2 instance that the UDDI Registry is using:
 

```
./home/db2inst1/sqllib/db2profile
```
6. Start the application server on which the UDDI Registry is to run. For example:
 

On Windows:

```
startserver server1
```

On Unix platforms:

```
./startServer.sh server1
```
7. Copy the *appserversetupuddi.jacl* script from the UDDIReg/scripts subdirectory of the deployment manager install tree to the WebSphere Application Server bin subdirectory (for example, on Windows, C:\Progra~1\WebSphere\AppServer\bin).
8. Change directory to the WebSphere Application Server bin subdirectory.
9. Create a JDBC driver and datasource to provide access to the UDDI20 database, and install the UDDI Registry application. To do this run the wsadmin tool with the script *appserversetupuddi.jacl* as input, on the target application server, using the following command syntax. (Either run this script from the UDDIReg/scripts subdirectory where it is located, or copy it to some other suitable directory. Note that the wsadmin tool is located in the WebSphere bin subdirectory.)

```
wsadmin -f appserversetupuddi.jacl
 uddi-ear-location
 servername
 nodename
 WebSphere-lib-subdirectory
 dbname
 db2userid
 db2pwd
 db2ziplocation
 > setupuddi.log
```

where

- *uddi-ear-location* is the fully-qualified path to the uddi.ear file in the installableApps subdirectory, specified **using forward slashes regardless of platform**.

For example, on Windows:

```
C:/Progra~1/WebSphere/AppServer/installableApps/uddi.ear
```

and on Unix platforms

```
/opt/WebSphere/AppServer/installableApps/uddi.ear
```

- *servername* is the name of the application server on which the UDDI registry is to run; for example: server1. Note that the name of the server is case sensitive.
- *nodename* is the name of the WebSphere node on which the application server, *servername*, is running. Typically, this will be the same as the machine name. Note that the name of the node is case sensitive. Typically this is the machine name.



- *WebSphere-lib-subdirectory* is the fully-qualified path to the WebSphere Application Server lib subdirectory, specified **using forward slashes regardless of platform**. For example:
  - On Windows: C:/Progra~1/WebSphere/AppServer/lib
  - On Unix: /opt/WebSphere/AppServer/lib
- *dbname* is the name of the UDDI Registry database under DB2. You should specify UDDI20 for this parameter

**Note:** If a remote DB2 system is being used the *dbname* must be the alias created when access to the remote database was set up on the local system.

- *db2userid* and *db2pwd* are a valid DB2 userid and password with administrative privileges, as specified in an earlier step.
- *db2ziplocation* is the path under which you have installed DB2, specified **using forward slashes regardless of platform..** For example, for Windows, this might be C:/Progra~1/SQLLIB/java/db2java.zip or, for Unix platforms it might be /opt/SQLLIB/java/db2java.zip.
- *> setupuddi.log* is an optional parameter to direct the output to a log file as opposed to the default (which is to the screen)

The appserversetupuddi.jacl completes the following actions:

- a. Creates a JDBC driver named *UDDI.JDBC.Driver.<nodeName>.<server>* and a datasource named *UDDI.Datasource.<nodeName>.<server>* (where *<nodeName>* is the name of the target node and *<server>* is the name of the target server), and will replace any existing driver and datasource of that name.
- b. Checks whether the WebSphere UDDI Registry application is already installed and, if so, stop the application and uninstall it.
- c. Installs the WebSphere UDDI Registry, then starts it.

**Note:** The setup script, appserversetupuddi.jacl, cannot be used to install the UDDI Registry application into a clustered application server. It is possible to cluster the UDDI Registry application by installing UDDI into an unclustered application server using the setup script, and then clustering that application server.

Return to the next step in the parent task Installing and Setting up a UDDI Registry.

---

## Reinstalling the UDDI Registry application

If you wish to reinstall the UDDI Registry, follow the appropriate section below.

### Reinstalling into a deployment manager cell

If you wish to reinstall the UDDI Registry into the target application server, for example because you wish to alter certain aspects of its configuration using AAT, rerun the setupuddi.jacl script (described in the appropriate link as follows):

- "Setting up the UDDI Registry to use Cloudscape within a deployment cell"
- "Setting up the UDDI Registry to use DB2 within a deployment cell"

**Note:** If you decide to change from using Cloudscape as your persistence store to DB2, or vice versa, first remove UDDI from the application server using removeuddi.jacl. You can then run setupuddi.jacl to reinstall UDDI with the new type of persistence store.

If you make such a change, then any data that you had previously stored is no longer be accessible.

### Reinstalling into a single appserver



Remove the UDDI Registry application in the same manner as any other Enterprise Application and then install as described in the appropriate link:

- Setting up the UDDI Registry to use Cloudscape in a single AppServer
- Setting up the UDDI Registry to use DB2 in a single AppServer

**Note:** If you decide to change from using Cloudscape as your persistence store to DB2, or vice versa, first remove UDDI from the application server using `appserverremoveuddi.jacl`. You can then run `appserversetupuddi.jacl` to reinstall UDDI with the new type of persistence store.

If you make such a change, then any data that you had previously stored is no longer accessible.

---

## Applying Service updates to the UDDI Registry in a Network Deployment and single application server environment

When applying service to your WebSphere and UDDI Registry, you need to be aware of the following.

If you are running in a Network Deployment cell, and you apply service to the UDDI Registry (on the deployment manager), the updates are automatically sent to all nodes running UDDI that are federated to that cell.

If, however, you are running in a stand-alone single application server (which is not known to the deployment manager), the updates are not automatically sent out. In this case, you must read the PTF Release Notes for instructions on how to do this.

In addition, there are some updates to files and scripts that may not automatically be sent to the application servers in either a Network Deployment or single application server environment. Please read the PTF Release Notes for details

---

## Using a remote DB2 database for UDDI

The WebSphere Application Server UDDI Registry can use a remote UDDI DB2 database installed on a DB2 system running on another machine.

It is possible to connect a DB2 system on one host to use a DB2 database on a remote host. Refer to the DB2 Information Center which supplies information on using the Client Configuration Assistant to achieve this. For the documentation below, the term *'local system'* refers to the system running the WebSphere Application Server on which the UDDI Registry is installed and the term *'remote system'* refers to another system on which the UDDI DB2 database exists. The UDDI DB2 database is created by running `SetupDB2UDDI.jar` on the remote system as described within "Setting up the UDDI Registry to use DB2 within a deployment manager cell" or "Setting up the UDDI Registry to use DB2 in a single appserver" and referred to in the last task in the list below. Refer to the DB2 specific version documentation for details.

Here is an example for DB2 Version 8:

1. On the remote host, define the UDDI Registry by entering: `java -jar SetupDB2UDDI.jar` and following the prompts. The `DB2userid` and `DB2password` specified here will be the ones that are used when `setupuddi.jacl` (or `appserversetupuddi.jacl`) are run further in these instructions.
2. Use DB2's Client Configuration Assistant. When you receive the message "Welcome to the Assistant - would you like to add a database now?" - select "yes"
3. On the next screen, click the radio button to "Manually configure a connection to a database" and click on "Next".
4. Click TCP/IP on the next screen, and select "The database physically resides on a host or OS/400 system. Click the radio button "Connect directly to the server" and then click "Next".

5. The next screen defines the TCP/IP host name and Port Number of the remote system. For the Port Number the use of 50000 is recommended. You are also asked for the service name (for example db2c\_db2inst1). Click "Next" when completed.
6. The next screen identifies the name of the UDDI database on the remote system and the alias name to be used on the local system. This alias can be anything suitable, including UDDI20 providing there is not a local database of that name.
7. The next screen shows "Register this database for ODBC". Leave this as is and click "Next".
8. The next screen asks for the Operating System (for example AIX or Windows) and the Remote Instance Name (for example db2inst1). Click "Next" when completed.
9. It is not necessary to go any further on this assistant. Click on "Finish" to define the remote database to DB2.
10. Another windows displays stating "Test the Connection". Enter the db2userid and db2password for the remote instance in the appropriate fields and click "Test Connection".
11. Click "Cancel" and the alias is now correctly set up.

**Note:** If the user is running a DB2 Client only, and has not got the DB2 Client Configuration Assistant available there are two options available:

- Obtain and install DB2 Administration Server (which contains the Client Configuration Assistant).
- Run the commands from the command line:
  1. Find the port number used by the remote instance you are trying to connect to:  
On the remote host as the DB2 instance user:
    - a. Enter the command: *"db2 get dbm cfg"*
    - b. Find the line containing *".....(SVCENAME) = <service\_name>"* Make a note of *<service\_name>*
    - c. There is a services file on each supported DB2 platform that changes location depending of the OS. For UNIX platforms this file might typically be found in */etc/services*, and on Windows this file might be found in *C:\WINNT\system32\drivers\etc\services*
    - d. Search this file and find the line for the *<service\_name>* discovered above. Use this port number for the catalog node command. For instance, the service may be *"db2c\_db2inst1"* and the port 50000, depending on the DB2 configuration.
  2. The following commands catalog a node and database alias (on the local host) to allow remote connectivity:
    - a. *db2 catalog tcpip node <nodename> remote <remote\_hostname> server <port>*  
Where *<nodename>* is a node identifier that you choose, *remote\_hostname* is the hostname of the remote host, and *port* is the port number discovered above
    - b. *db2 catalog db UDDI20 at node <nodename>*  
Where *UDDI20* is the name of the UDDI Registry database on the remote system and *<nodename>* is the identifier you chose in the previous step.
  3. You can now to the remote database. Test this by:
    - a. *db2 connect to uddi20 user <db2userid> using <db2password>*  
Where *<db2userid>* is the userid of the DB2 instance on the remote host and *<db2password>* is the password for the DB2 instance on the remote host.

Run the *setupuddi.jacl* (or *appserversetupuddi.jacl*) command on the local system as described in "Setting up the UDDI Registry to use DB2 within a deployment manager cell" or "Setting up the UDDI Registry to use DB2 in a single appserver", ensuring that you specify the alias on the local system as the UDDI database parameter, and the DB2 userid and password on the remote host as the userid and password parameters.

## IP Addressing

If you are experiencing problems accessing the remote host, one of the possible causes might be IP addressing. Refer to the "UDDI Troubleshooting Tips" section for a possible solution.

The remote database is now ready for use. Continue with your next task.

---

## Removing the UDDI Registry application from a deployment manager cell

To completely remove the UDDI Registry application from the target application server in the deployment manager cell, run the *wsadmin* (*wsadmin.sh* on Unix Platforms) script *removeuddi.jacl*, which is located in the UDDIReg/scripts directory of the deployment manager install tree.

If the target server specified on invoking *removeuddi.jacl* is running at the same time, the script stops the server and restarts the server when it has completed its operations.

At a command prompt enter:

```
wsadmin -f removeuddi.jacl
 servername
 nodename
 > removeuddi.log
```

Where *servername* and *nodename* are the server and node where you have deployed the UDDI Registry application. By default output will go to the screen, but, optionally, you can specify '> *removeuddi.log*' to direct output to a log file. For example,

```
wsadmin -f removeuddi.jacl server1 myriad
```

will remove the UDDI Registry application and related files from server *server1* running in node *myriad*, and will send any messages to the screen.

---

## Removing the UDDI Registry application from a single application server

To completely remove the UDDI Registry application from a stand-alone application server run the *wsadmin* script *appserverremoveuddi.jacl*, which was installed into the UDDIReg/scripts directory when you installed the UDDI Registry as part of a Network Deployment install.

At a command prompt enter:

```
wsadmin -f appserverremoveuddi.jacl
 servername
 nodename
 > removeuddi.log
```

where

- *servername* and *nodename* are the name of the stand-alone application node in which it runs (these are the names that you specify when you run *appserversetupuddi.jacl* to install the UDDI Registry application).
- by default output will go to the screen, but, optionally, you can specify '> *removeuddi.log*' to direct the output to a log file.

For example:

```
wsadmin -f appserverremoveuddi.jacl server1 monolith
```

will remove the UDDI Registry application and related files from server *server1* running in node *monolith*, and will send any messages to the screen.

---

## Configuring the UDDI Registry

The UDDI Registry is supplied as a J2EE application file, `uddi.ear`. This is installed into the WebSphere Application Server during installation. If you want to change any of its configuration properties using Assembly Toolkit see [Configuring SOAP properties with the Assembly Toolkit](#).

You can configure the following aspects of the UDDI Registry:

- Configuring global UDDI properties
- Modifying the database userid and password
- Configuring security properties
- Configuring the UDDI User Console (GUI) for multiple language encoding support
- Customizing the UDDI User Console (GUI)
- Configuring SOAP interface properties
- **5.1+** Configuring SOAP properties with the Assembly Toolkit
- Configuring SOAP properties in the deployment descriptor

### Configuring global UDDI properties

To modify any of the global UDDI properties, edit the file called `uddi.properties`. More than one version of this file exists and the version you need to edit depends on:

- whether you are in the installation phase or are updating the properties as a post installation step
- whether you are configured for a deployment manager or base application server environment

The location of the file you should edit will be one of the following:

#### **Deployment Manager Configurations**

1. If you are in the process of installing the UDDI Registry application for the first time into a deployment manager cell and wish to make some generic changes before deploying it in the cell, the `uddi.properties` file will be located in the `<DeploymentManager-install-dir>/properties` directory. If you are reinstalling the UDDI Registry application into a deployment manager cell, then you should edit the file in the location described in step 2.

**Note:** In a deployment manager configuration some properties (such as `persist` and `getServletURLprefix`) are dynamically set up in the `uddi.properties` file, during subsequent installation processing.

2. If the UDDI Registry is already configured into an application server within a Deployment Manager cell (that is you are undertaking post installation configuration changes), the `uddi.properties` file you should edit is located in the configuration repository, under the deployment manager filing system; that is in `<DeploymentManager-install-dir>config/cells/<cellname>/nodes/<nodename>/servers/<servername>`, where `<cellname>` is the name of the deployment manager cell, `<nodename>` is the name of the node in which the application server is installed, and `<servername>` is the name of the application server in which you have installed the UDDI Registry.

#### **Application Server Configurations**

1. If you are in the process of installing the UDDI Registry application into an application server only environment you will be advised during the installation process when to make changes to the `uddi.properties` file.

**Note:** In contrast with the deployment manager configuration, UDDI properties are not dynamically set during installation processing.

2. If the UDDI Registry is already configured into a single application server that is **not** part of a deployment manager cell (i.e. you are undertaking post installation configuration changes), then the

*uddi.properties* file will be located in the properties subdirectory of the WebSphere Application Server in which you have installed the UDDI Registry application, that is <ApplicationServer-install-dir>/properties directory.

The properties that can be changed within *uddi.properties* are as follows:

- The `dbMaxResultCount`, which is the limit on the number of rows of information that should be returned on Find requests, and will apply if the request does not specify a `maxRows` limit itself (or if it specifies a limit that exceeds this value). The initial value for this in *uddi.properties* is 100.
- The `persist`, which indicates what database is to be used as the persistence store for the UDDI Registry database. If you have installed the UDDI Registry into an application server within a deployment manager cell, then the `persist` property will have been set to the correct value for you. If you change this value, you must also ensure that you have a UDDI Registry database created using the chosen database product (for more details about the UDDI Registry database, refer to the section on "Installing the UDDI Registry"). You should also be aware that any data published to the UDDI Registry with one setting of the `persist` property will **not** be accessible when running the UDDI Registry application with a different setting for the `persist` property. The valid values for the `persist` property are:
  - `persist=DB2`  
indicating that DB2 is to be used as the persistence store
  - `persist=Cloudscape`  
indicating that Cloudscape is to be used as the persistence store

The initial value for this in *uddi.properties* is Cloudscape.

**Note:** This property is dynamically set by the `setupuddi.jacl` script when installing into a deployment manager cell so in this case you should not need to modify it.

- The default language to be used on a publish request as the `xml:lang` attribute when one is not specified. The initial value for this in *uddi.properties* is en-US. This property must contain one of the valid `xml:lang` values.
- The UDDI site operator name. This is a string that is stored in every registry object, to indicate the operator of the UDDI Registry. The initial value for this in *uddi.properties* is `www.mycompany.com/uddi`. This property does not have any particular functional use, so its value can be set to any string that you feel is suitable.
- The maximum number of search keys that can be used on find API requests. The initial value for this in *uddi.properties* is 5.
- The `getServletURLprefix` and `getServletname` name, used to build up the discovery URL. The initial values for these are `http://localhost:9080/uddisoap/` and `get`. If you have installed the UDDI Registry into an application server within a deployment manager cell, then the `getServletURLPrefix` property will have been set for you using the value you specified as a parameter to the setup script. You are recommended to set suitable values for these properties before you first use the UDDI Registry.

**Note:** This property is dynamically set by the `setupuddi.jacl` script when installing into a deployment manager cell so in this case you should not need to modify it.

### **Applying these changes to your system**

For your changes to take effect, you must do one of the following:

- If you are in the process of installing the UDDI Registry application for the first time, return to the original topic and complete the installation steps. Any changes you have made are picked up during this subsequent processing.
- If you have made post installation changes in a base application server only environment, you should stop and restart the UDDI Registry application using the WebSphere administrative console.
- If you have made post installation changes in a deployment manager environment you should:
  1. Run a Full Resynchronization for the node where the UDDI Registry runs. This can be done from the WebSphere Network Deployment Administrative Console under section Systems Administration ==> Nodes. Select your node, and then click 'Full Resynchronization'.

- It is important that you do a 'Full Resynchronization' and not just a 'Synchronize'.
2. Stop and restart the UDDI Registry application using the WebSphere administrative console

## Modifying the database userid and password

If you use DB2 as the persistence store for the UDDI Registry, and you need to change the database userid and/or password, alter the user and password values in the custom properties of the 'UDDI Datasource', which can be edited from the WebSphere administrative console. The UDDI.Datasource is under WAS40 datasources within the UDDI.JDBC.Driver, which is itself found under JDBC Providers under Resources. Do not alter the databaseName.

## Configuring security roles

Each interface to the UDDI Registry (either through SOAP, EJB or the GUI) is supplied with two roles:

### Publish role

mapped to AllAuthenticatedUsers. By default, this is configured to use SSL (that is HTTPS), but this only applies when WebSphere security is enabled.

### Inquiry role

mapped to Everyone. By default, this is configured to use HTTP (that is not SSL).

The security roles and use of SSL can be altered by users through the Administration Console.

Authentication uses the standard WebSphere facilities and there is no separate registration function for the Registry. If WebSphere security is enabled, the you will need to supply your WebSphere userid and password for Publish functions (unless you have changed the supplied Publish role).

You will need to set up WebSphere security configuration to be used by UDDI. For information on achieving this, refer to "Configuring Secure Sockets Layer" in the Information Center. It is expected that, for development use, security will be disabled and security will be enabled for production environments.

The SOAP interface also supports the UDDI API for get\_authToken and discard\_authToken API but use of this is optional.

- If security is disabled and get\_authToken is not called, the default user, UNAUTHENTICATED, is used.
- If security is disabled and get\_authToken is called, the specified userid is used (but the password is not checked).
- If WebSphere security is enabled, it takes priority over UDDI authentication, but if the Publish role is mapped to Everyone, get\_authToken must be used and the userid and password will be checked by WebSphere.

The Security Roles provided with the UDDI Registry are as follows:

- GUI\_Publish\_User
- GUI\_Inquiry\_User
- SOAP\_Publish\_User
- SOAP\_Inquiry\_user
- EJB\_Inquiry\_Role
- EJB\_Publish\_Role



## Configuring the UDDI User Console (GUI) for multiple language encoding support

If you want to use multiple language encoding support in the User Console (GUI), you must configure the application server into which the UDDI Registry application is installed with UTF-8 encoding enabled. To do this, refer to "Configuring application servers for UTF-8 encoding" in the WebSphere Information Center.

## Customizing the UDDI User Console (GUI)

The look and feel of the UDDI console is determined by the styles defined in the *uddi\_gui.css* file which is located in the `/gui.war/theme` directory of the installed UDDI Registry application directory. The UDDI Registry application directory will be one of the following, depending on where you have installed the UDDI Registry:

- If you have installed the UDDI Registry into an application server within a deployment manager cell, the directory is `<AppServer-install-dir>/installedApps/<currentcell>/UDDIRegistry.<nodename>.<servername>.ear/gui.war`
- If you have installed the UDDI Registry into a single application server which is not part of a deployment manager cell, the directory is `UDDIRgistry.ear` under the `installedApps` directory of the WebSphere Application Server in which you have installed the UDDI registry application as shown in the example below.

The contents of this file can be edited to change the colors, fonts and font sizes according to the user's preference.

The content and layout of the UDDI User Console is provided by Java Server Pages (JSP), which can be customized by a programmer who is familiar with JSPs. The JSP pages are found in the `uddi.ear` enterprise application, which is under the `installedApps` subdirectory of the WebSphere AppServer installation. To locate the JSPs, expand the `UDDI_Registry.ear`, open the `gui.war`, and they are located under `WEB-INF` in the `pages` subdirectory. So, on a Windows system that has WebSphere installed in the default location, the JSP files will be found in

```
<AppServer-install-dir>\installedApps\<nodename>\UDDI_Registry.ear\gui.war\WEB-INF\pages
```

These JSP pages also contain some application logic (as opposed to presentation logic) that should not be changed.

## Configuring SOAP interface properties

You can configure the following SOAP interface properties:

- *defaultPoolSize* - the number of SOAP parsers with which to initialize the parser pool for the SOAP interface. You can set this independently for the Publish (`uddipublish`) and Inquiry (`uddi`) APIs. For example, if you expect more inquiries than publish requests via the SOAP interface, you can set a larger pool size for the Inquiry API. The default initial size for both APIs is 10.
- The *context root* used for the Publish and Inquiry APIs, which forms a part of the URL by which they are accessed. By default this is `/uddisoap`.
- Whether the API is to be secure (via HTTPS) or insecure (via HTTP). The default is to use HTTPS.

To configure the following SOAP interface properties, use either of the following methods:

- Configuring SOAP properties with the Assembly Toolkit (the recommended option, especially for a production environment)
- Configuring SOAP properties in the deployment descriptor for the SOAP module in the UDDI application directly. This option is faster and may be the preferred method in a test environment.



## Configuring SOAP properties with the WebSphere Assembly Toolkit

To configure SOAP properties by using the WebSphere Assembly Toolkit, complete the following steps:

- Select *Update* and click on the Application icon.
- Select the *uddi.ear* file (this is placed, by the UDDI installation, into the UDDI install directory (e.g. *C:\WebSphere\installableApps\uddi.ear*).
- Expand the *uddi.ear* icon on the left-hand pane in the AAT.
- Expand the *Web Modules* tree.
- Expand the *uddi Soap* tree
- To change the *defaultPoolSize*, expand Web Components and then *uddipublish* (for the publish API) or *uddi* (for the inquiry API).
  - Click on *Initialization Parameters* which will show the *defaultPoolSize* parameter in the upper right-hand pane. This can be edited in the lower right-hand pane.
- To change the *context root*, click on *UDDI Soap* which displays general information about the SOAP module in the lower right-hand pane in Assembly Toolkit. The *context root* can be edited in this pane.
- To change the publish API to use HTTP (instead of HTTPS), click on *Security Constraints* and change the *Transport Guarantee* from Confidential to none.
- Having made any changes above, you must now save them. To do this, click on *File -> Save (or Save As)* to save your changes.
- Redeploy the *uddi.ear* to WebSphere, by first removing it and reinstalling it via the Administrator's Console.

## Configuring SOAP properties in an application that is already deployed

To configure SOAP properties after the UDDI application has been installed:

1. Edit the deployment descriptor for the SOAP module (*web.xml*). This file is located in the *WEB-INF* subdirectory of the *uddi.ear* application in the installed applications within the WebSphere install directory (for example, *<WebSphere-install-dir>\installableApps\uddi.ear\soap.war\WEB-INF*).
2. Stop and restart the application server for the changes to take effect.

---

## Administering the UDDI Registry

Perform the following tasks to administer the UDDI Registry:

- Running the UDDI Registry
- Backing up and restoring the UDDI Registry database

## Running the UDDI Registry

### Starting the WebSphere Application Server in which the UDDI Registry is deployed

After a reboot, or at any time required, the server can be started by running *startServer* (on Windows) or *startServer.sh* (on Unix and Linux platforms).

On Unix and Linux platforms run the *db2profile* script before issuing the *startServer.sh server1* command. This script is located within the DB2 instance home directory under *sqllib* and is invoked by typing:

```
. /home/db2inst1/sqllib/db2profile
```

**Note:** In the above example, notice that the *'.'* is followed by a single space character.

**Note:** On Unix and Linux platforms the DB2 user **must** have a *db2profile* at *\$HOME/sqllib/db2profile*

By default, the UDDI Registry is started automatically when the application server is started. In order to stop and restart it, use the administrative console.

You can also use the administrative console to change this default behavior.

## Backing up and restoring the UDDI Registry database

If you want to protect the data in your UDDI Registry database, you can back up and restore the database using the facilities of the database product. For DB2, you can do this by using the export and import utilities of the DB2 Control Center. For Cloudscape you can simply use operating system tools to copy the database directory. Refer to the database product information for more details.

DB2 allows for dynamic backup, but, if you are using static backup, stop any UDDIReg applications beforehand.

The UDDI Registry database is called UDDI20, and the tables that should be backed up are:

- ADDRESS
- ADDRLINE
- BSERVICE
- BTEMPLATE
- BUSINESS
- CATEGORY
- CATEGORYBAG
- CONTACT
- DESCR
- DISCOVERYURL
- EMAIL
- EXTCATEGORY
- IDENTIFIERBAG
- INSTANCEDetail
- NAMEELEMENT
- OVERVIEWDOC
- PHONE
- PUBLISHERASSERTION
- SERVICEPROJECTION
- TMODEL
- VALIDATIONCACHE
- VALIDATIONSERVICES

---

## UDDI4J programming interface

The following considerations are specific to the support for UDDI4J specification provided by WebSphere Application Server:

- **UDDI4J class libraries provided.**

WebSphere Application Server provides two UDDI4J class libraries:

**uddi4jv2.jar**

This class library contains classes which support Version 2 of the UDDI specification.

**uddi4j.jar**

This class library is provided for compatibility with WebSphere Application Server and supports Version 1 of the UDDI specification. The classes in this library are deprecated.

- **UDDI4J error handling**

When invoking UDDIProxy inquiry methods, UDDIException is thrown if errors are received from the UDDI proxy. UDDIException can contain a DispositionReport with information about the error. APIs that do not return a data object, provide the disposition report. SOAPException is thrown if a communication error occurs or if the resulting data cannot be parsed as a valid SOAP message.

See the article [Web services: Resources for Learning](#), to learn more about the use of SOAP, Universal Description, Discovery and Integration (UDDI) Registry, UDDI4J, programming specifications and examples.

---

## UDDI user console

This topic describes the layout of the UDDI user console (also referred to as the Graphical User Interface (GUI)), which you can use to interact with the IBM WebSphere UDDI Registry.

For information about how to display the UDDI user console, see [Displaying the user console](#).

If you will be using the UDDI console, it is recommended that you configure the application server into which you have installed the UDDI Registry for UTF-8 encoding support: see ["Configuring the UDDI User Console for multiple language encoding support"](#).

- The user console provides a graphical user interface to the majority of the UDDI Version 2 API. It is not intended to support the full API set: there is some focus on inquiry operations, as the main purpose of the UDDI user console is to allow users to issue inquiry requests and to familiarize themselves with general UDDI concepts. This section documents those areas for which support through the user console is not provided, together with other known restrictions to the user console.
  - General
    - Help is provided in the form of explanatory text on the screens.
    - Maximum rows cannot be specified on finds. The single maximum rows value for the registry can be set through the *dbMaxResultCount* global configuration property. For more information on setting this property see [Configuring global UDDI properties](#)
  - Find business
    - The business identifier feature is not supported.
  - Find service type
    - The business identifier feature is not supported.
  - Add business
    - You must supply the business contact as a name and role (no other information is supported).
  - Add service type
    - You can enter the overview URL, but only with one description in English.
  - Add service
    - There is no support for entering a Hosting Redirector, nor for adding an overviewURL.
- **Note:** The UDDI Version 2 specification states that when a tModel is deleted, it should not be physically deleted. This allows the tModel to be reinstated. One effect of this is that, if you delete a tModel using the UDDI user console, the tModel is still visible through the Show Owned Entities display.

The UDDI user console is split into three distinct areas. At the top of the screen are buttons that activate various functions in the areas below this bar. These buttons are:

**Home** Returns you to the IBM WebSphere UDDI Registry welcome page

**Find** Activates the Find tab on the frame below to the left

**Publish**

Similarly activates the Publish tab on the frame below to the left

Below the WebSphere UDDI Registry banner the screen is split into two parts. On the left are the two tabs mentioned above, the Find and Publish tabs.

### Find tab

The Find tab is in two parts. At the top, a **Quick Find** service is provided. There are three radio buttons to enable a choice of 'service', 'business' and 'technical model' finds. Below these radio buttons is a text entry box for entering the name to search for and, beneath this, a 'Find' link to start the search. Comments are provided to show the user the wildcard character. The results of clicking on the 'Find' link are shown in the detail frame to the right.

Beneath the Quick Find is a section for **Advanced Find** functions which enables the user to choose which entity they want to perform an advanced search on. There are three links: Find services, Find businesses and Find technical models. Clicking one of these links displays the corresponding advanced search form in the frame to the right, which the user may use to enter search criteria. To initiate a Find, the user must first enter a search path (the % wildcard may be used) and then click the blue Add link to enter the search. Then click on the 'Find Services' (or 'Find Businesses/Find technical models) link below to initiate the Find operation. The **Locator** section has a link (marked in blue with the words "**Show category tree**") which displays the tree from which the user can select categories (or taxonomies). This is shown in the left-hand frame. In the advanced search form there are two links to start the search (mid-way down and at the bottom).

The results of clicking either of the two links to start the search are displayed in the same detail frame.

## **Publish tab**

The *Publish* link on the top banner activates the Publish tab in the navigation frame to the left. The Publish tab is split into three distinct sections.

### **1. Quick Publish Function**

The top part is a **Quick Publish** section to allow the user to publish a business or technical model by name only. There are two radio buttons to enable a choice of 'business' or 'technical model'. Below these radio buttons is a text entry box for entering the name to assign to the selected entity and, beneath this, a blue 'Publish now' link to publish the entity. The results of clicking on the **Publish now** link are shown in the detail frame to the right.

### **2. Advanced Publish Functions**

To publish an entity with more detail, such as with multiple names, descriptions and categories, use the **Advanced Publish** section below this. The comments below each link ('Add a business' and 'Add a technical model') describe individual functions. Clicking one of these links displays the corresponding advanced publish form in the detail frame where the user may enter details about the entity they want to publish. As in the Advanced Find functions described above, there are two links to publish a business or technical model (one towards the top of the form and the other at the bottom). Similarly the **Locator** section allows taxonomies to be shown in the left frame from which the user can select categories.

Following entry of the relevant details on the **Advanced Publish** section, the user must click on the **Publish Business** bar in order for the business to be published to the UDDI Registry.

### **3. Registered Information**

Below the Advanced Publish section is a **Registered Information** section which has a link to **Show Owned Entities** in order to show the businesses, services and technical models registered to the individual user, and pending business relationships. Clicking the **Show Owned Entities** link displays the **Show Owned Entities** page in the detail frame at the right. The **Show Owned Entities** page is organized in three sections: **Registered Businesses**, **Pending Business Relationships** and **Registered Technical Models**. Each section shows the number of registered items.

#### **Edit and Delete Businesses**

Users can **Edit** or **Delete** businesses owned by them by clicking the appropriate links in the **Actions** column.

After an **Edit** or **Delete** function has been completed, the user **must** click on the **Update Business** bar in order to publish the changes to the UDDI Registry.

After Deleting a Business the user must confirm the deletion by clicking on the 'Delete this Business' link.

#### **Adding a Service to a Business**

Services are added to a business by clicking the **Add a Service** link in the **Services** column of the **Registered Businesses** section.

After the **Add a Service** function is complete, users **must** click on the **Publish Service** bar in order to publish the service to the UDDI Registry.

## Referencing a Service from a Business

Services can also be 'referenced' by a business as if the business was the owner of the service. This 'service projection' is performed by clicking the **Reference a service** link in the **Services** column. Services associated with a business, whether they are owned or referenced, can be displayed by clicking the **Show services** link. This acts as a toggle between displaying services available for editing or deleting, and hiding them.

## Adding a relationship to another Business

A business can be associated with another business in the UDDI Registry and this function is performed by clicking the **Add a relationship** link in the **Actions** column of the Registered Businesses section. Clicking the **Show related businesses** link in the **Actions** column displays a list of any completed business relationships.

The **Pending Business Relationships** section shows all incomplete publisher assertions, where only one party has asserted a relationship and is waiting for the other party to make the same assertion. This section reminds the user of any relationships that involve their businesses. Once both parties have asserted the same relationship between two businesses, the relationship moves from the **Pending Business Relationships** section and appears in the list of relationships displayed after clicking the **Show related businesses** link in the **Registered Businesses** section.

## Technical Models

Technical Models owned by the user are shown in the bottom **Registered Technical Models** section. As for businesses, users can Edit or Delete technical models owned by them by clicking the appropriate links in the **Actions** column.

**Note:** Users should take note that deletion of Technical Models (tModels) does **not** cause them to be physically deleted, but hidden. This is in accordance with the UDDI Registry V2.0 specifications. After deletion Technical Models are shown under the "**Shown Owned Entities**" link on the publish page but not via the Find links on the Find page. ALL other entities are deleted from the UDDI Registry in the normal way.

## Example of publishing a Business, Service and tModel with the User Console

For the example, here, we will assume a business called Mondeo Cars that sells used cars

### 1. Add the Business

Click on the Publish tab in the left hand navigation frame. Then click on 'Add a business' in the Advanced Publish in the left pane. This takes you to a 'Publish Business' pane on the right. Start by adding your Business Name in the text field labelled (Mondeo Cars in this example) and select a language and then click on the blue Add link to the right. This adds the business name (but the business is not yet published - more about which is explained later). Below the Business Name is an area called Descriptions - it allows free text to be added to describe the business - if you enter anything here you must click on the blue Add link to the right to insert the description.

The next section/area is the Locator area which can be used to describe the business according to what categories it falls into. This example uses a Used car dealership. Within the NAICS taxonomy (which you may view by clicking on 'Show category tree' and then expanding NAICS) this is a Retail Trade [44] entry which, on expansion, has Motor Vehicle and Parts Dealers [441] and, again on expansion, has automobile Dealers [4411] and Used Car Dealers [44112]. This fits the Business perfectly, so clicking on Used Car Dealers will enter the Key Name and Key Value into the business. For Checked Categorizations (such as NAICS) the Key name is not checked but the Key Value is checked. It should be noticed that for unchecked categorizations (such as 'other' or unspsc') the Key value is not checked either. If the locator field has been added, then the blue Add link must be clicked. The final area is Contacts, which can have names and role information added if required. Again, the blue Add link must be clicked after adding the relevant information.

Once all the fields are filled in to the required level, the final action is to click on the Publish Business at the bottom of the form or at the top. This causes the business to be published to the UDDI Registry and a page is displayed showing the business details.

### 2. Add a Service

From the Publish tab, there is a 'Show owned entities' link. This shows the businesses owned by the current user in the Registry and the language to be used for a particular user. For Mondeo Cars, the user will see a 'Add service' button. Clicking this button shows the Publish Service form. The top part of the form is the Service Name field. After adding this name, the user must click on the blue Add link to enter the name. As in the Adding the Business form, each subsequent part must end with the blue Add link being clicked to add that part of the information to the service. The sections are (from top to bottom, Description (a free text area), Access Points (to add link points to the Service), Locator (to add references to taxonomies to the service), and Technical Models (to associate existing tModels to the Service). After completion of those areas required, clicking on the 'Publish Service' button will Publish the service to the UDDI Registry with the current form contents.

### 3. **Adding a new technical model**

Clicking on the 'Add a technical model' link in the left frame opens up the Publish Technical Model form on the right. A tModel can only have one name - hence the lack of a blue Add link next to the Technical Model Name field. Beneath this field are other fields - Description (a free text area to describe the technical model), Locator (to describe the technical model with taxonomies, and an Overview URL (which gives a URL pointing to an overview document, a description of the document and a Language field). For each of these fields there is a blue Add link which must be clicked to add the relevant data. At the bottom of the form is a 'Publish Technical Model' link which will create the technical model in the UDDI Registry.

There is a Publish link at the top of the frame in each case also - after the Name section.

- Displaying the user console

## Displaying the user console

### Access without authentication enabled

This topic describes how to display the UDDI Registry user console (also referred to as the GUI). By default two URLs are supplied, one for inquiry (non-SSL) and one for publish (via SSL). This section describes the default behavior.

- For inquiry you can access the UDDI User Console by using the following URL in your Web browser:

`http://<hostname>:9080/uddigui`

- For publish you can access the UDDI user console by using the following URL in your Web browser:

`https://<hostname>:9443/uddigui`

**Note:** With WebSphere security disabled, all the publish operations are performed using a userid of UNAUTHENTICATED. Also, if you select a publish action on the GUI you will automatically be redirected to the SSL port.

### Access with authentication enabled.

If you have WebSphere security enabled, you can access the UDDI User Console through the two URLs as above, however, for publish request, you will be prompted for a WebSphere uid and password.

The user console displays the default frameset containing the header frame, navigation frame showing find options, and details frame. When you click the link to show the publish options in the navigation frame, you are challenged for a userid and password.

If WebSphere security is enabled and you try to access a publish action via an unsecured link, e.g. clicking the publish link on the navigation frame where the user console was opened with:

`http://<hostname>:9080/uddigui`

you are redirected to a secure logon screen. Inquire functions work as expected.



The uDDI Registry supports a number of security roles, including two for the user console. See [Configuring Security Roles](#) within this Information Center for more details on this topic.

## Custom Taxonomy Support in the UDDI Registry

The IBM WebSphere UDDI Registry is supplied with six published taxonomies (or categorization schemes) in the taxonomy data. Taxonomies can be either checked or unchecked, and this is indicated via a keyedReference in the categoryBag of the tModel that represents a taxonomy (a "categorization tModel"). These keyedReferences have the tModel key for uddi-org:types and are added to the categoryBag to further describe the behavior of the categorization tModel, as follows:

### checked

Marking a tModel with this classification asserts that it represents a categorization, identifier, or namespace tModel that has a validation service to check that category values are present in a specified value set.

### unchecked

Marking a tModel with this classification asserts that it represents a categorization, identifier, or namespace tModel that does not have a validation service.

Of these six published taxonomies, four are checked.

In the IBM WebSphere UDDI Registry and also in the IBM UDDI Business Registry (UBR), the validation of categories in checked taxonomies is performed against locally managed taxonomy data. The published taxonomies are:

Taxonomy name	Checked	Description	tModel key
ntis-gov:naics:1997	Yes	Business Taxonomy: NAICS (1997 Release)	uuid:C0B9FE13-179F-413D-8A5B-5004DB8E5BB2
uddi-org:iso-ch:3166-1999	Yes	ISO 3166-1:1997 and 3166-2:1998. Codes for names of countries and their subdivisions. Part 1: Country codes. Part 2: Country subdivision codes. Update newsletters include ISO 3166-1 V-1 (1998-02-05), V-2 (1999-10-01), ISO 3166-2 I-1 (1998)	uuid:4E49A8D6-D5A2-4FC2-93A0-0411D8D19E88
unspsc-org:unspsc	Yes	Product Taxonomy: UNSPSC	uuid:CD153257-086A-4237-B336-6BDCBDCC6634
unspsc-org:unspsc:3-1	No	Product Taxonomy: UNSPSC (Version 3.1)	uuid:DB77450D-9FA8-45D4-A7BC-04411D14E384
uddi-org:types	Yes	UDDI Type Taxonomy	uuid:C1ACF26D-9672-4404-9D70-39B756E62AB4
uddi-org:general_keywords	No	Special taxonomy consisting of namespace identifiers and the keywords associated with the namespaces	uuid:A035A07C-F362-44DD-8F95-E2B134BF43B4



Taxonomy data is provided in the IBM WebSphere UDDI Registry for all the above taxonomies, apart from the general keywords taxonomy (which is unchecked). The UDDI User Console (GUI) provided with the IBM WebSphere UDDI Registry uses a shortened label for taxonomies when displayed in the taxonomy tree view, or in a pull-down list of available taxonomies as follows:

Taxonomy Name (published)	Taxonomy name (as displayed in the UDDI user console)
ntis-gov:naics:1997	naics
uddi-org:iso-ch:3166-1999	geo
unspsc-org:unspsc	unspsc7
unspsc-org:unspsc:3-1	unspsc
uddi-org:types	udditype
uddi-org:general_keywords	other

This release of IBM WebSphere UDDI Registry (included with IBM WebSphere Application Server, Version 5.0.2) introduces the ability to add user-defined taxonomies, with available allowed values presented in the existing GUI taxonomy tree display. IBM WebSphere Studio Application Developer, Version 5.1 has a Web Services Explorer user interface that also allows addition and display of custom checked taxonomies. The publisher of a custom taxonomy's categorization tModel may specify a 'display name' for use in GUI implementations.

### Procedure for adding a Custom taxonomy

To add a custom taxonomy to the IBM WebSphere UDDI Registry requires you to perform two tasks: load the custom taxonomy data and publish a categorization tModel. Only when both are complete will the checked taxonomy be of practical use. Taxonomy data must be provided for validating checked taxonomies.

Taxonomy data *may* also be used by GUIs for unchecked taxonomies, but it is not a requirement and is usually only used for presentation of deprecated taxonomies, such as unspsc-org:unspsc.

If the taxonomy is checked, then any publish requests that have a categoryBag containing keyedReferences with the new categorization tModel will be validated. If there is taxonomy data corresponding to the categorization tModel in the registry database then only valid values will be accepted. If there is no taxonomy data in the database then **all** values will be rejected, and the publish request will fail. If the categorization tModel is unchecked, all values will be allowed, regardless of whether there is corresponding taxonomy data present in the UDDI Registry database.

### Suggested approach

The suggested way of introducing a new taxonomy is to:

1. Load custom taxonomy data into the UDDI Registry database using the UDDITaxonomyTools.jar utility (described below)
2. Publish the categorization tModel with a keyedReference of type 'general keywords' with keyname of 'customTaxonomy:key' and a keyValue matching the taxonomy name in the taxonomy data file (described below also)

**Note:** the SOAP and EJB interfaces will be able to make use of categorization tModels as soon as they are published. However, the UDDI Registry GUI will currently require a restart of the UDDI application because it currently gathers its list of categorizations for use in the taxonomy tree display when the application starts.

## Loading Custom Taxonomy Data

### Custom Taxonomy Data File Format

Taxonomy data is identified by a common taxonomy name, a unique code value, an optional description and a parent code that specifies its relationship with other code values. Taxonomy data must adhere to this format:

Column name	Maximum length	Description of use
<b>name</b>	8	Uniquely identifies the taxonomy within the registry
<b>code</b>	32	Unique value within the taxonomy used for validation
<b>description</b>	128	Typically used by GUIs and optionally in the keyedReference as the keyName value
<b>parentcode</b>	32	Indicates which existing <b>code</b> is the logical parent of this one, and is used in tree displays

Typically columns are delimited in the taxonomy data file by '#' characters as in this example:

```
food#00#Food#00
food#10#Fruit#00
food#101#Apples#10
food#102#Oranges#10
food#103#Pears#10
food#1031#Anjou#103
food#1032#Conference#103
food#1033#Bosc#103
food#104#Pomegranates#10
food#20#Vegetables#00
food#201#Carrots#20
food#202#Potatoes#20
food#203#Peas#20
food#204#Sprouts#20
```

In the example, 'Food' is the description for the root node with child nodes of 'Fruit' and 'Vegetables' (both of these have parentcode values the same as the code value for 'Food').

The taxonomy data in the example file could then be rendered in a tree like this:

```
Food
 Fruit
 Apples
 Oranges
 Pears
 Anjou
 Conference
 Bosc
 Pomegranates
 Vegetables
 Carrots
 Potatoes
 Peas
 Sprouts
```

The file must be saved in UTF-8 format.

The following taxonomy names are reserved within the IBM WebSphere UDDI Registry and should not be used for custom taxonomy files: **naics**, **geo**, **unspsc**, **unspsc7**, **other**, **udditype**. Any attempts to publish a categorization tModel using these values for a customTaxonomy:key are rejected. If these names are used in custom data files and the data is imported it is indistinguishable from taxonomy data with the same name.

## UDDITaxonomyTools.jar

A utility is provided to load taxonomy data into the IBM WebSphere UDDI Registry, rename existing taxonomy data and remove existing taxonomy data, for both IBM DB2 and Cloudscape databases. The usage for each database and platform is identical:

Usage: java -jar UDDITaxonomyTools.jar {function} [options]

function:

```
-load <path> Load taxonomy data from specified file
-rename <old> <new> Rename existing taxonomy
-unload <name> Unload existing taxonomy
```

options:

```
-properties <path> Specify location of configuration file
```

**Note:** Ensure that the command window from which the UDDITaxonomyTools.jar is run is using a suitable codepage and font for displaying the characters contained in the taxonomy name.

Use of an incorrect codepage/font may result in unclear messages on a successful load, and create difficulty using the -unload and -rename options.

UDDITaxonomyTools.jar is located, by default, in the <DeploymentManager-install-dir>/UDDIReg/scripts directory.

The following section explains in more detail how to use the utility's commands and parameters. The configuration file, if specified by the optional properties parameter, determines the database driver, authentication information and delimiters. The contents are as follows (typical data for DB2 installation shown):

Property and example data (for DB2)	Comments
classpath= "c:/program files/sqllib/java12/db2java.zip; c:/tools/UDDITaxonomyTools.jar"	Classpath including database driver <i>and</i> the UDDITaxonomyTools.jar*
database.driver.className=COM.ibm.db2.jdbc.app.DB2Driver	Fully qualified classname of the database driver class
database.url=jdbc:db2:UDDI20	JDBC URL of the database
database.userName=db2admin	Database userid (DB2 only)
database.password=db2admin	Database password (DB2 only)
column.delimiter=#	Column delimiter used in taxonomy data files
string.delimiter=\"	Field delimiter (must be different to the column.delimiter value)

\* the classpath needs to be enclosed in quotes if the path includes space characters. Also, the UDDITaxonomyTools.jar filepath itself must be appended to the classpath (if the working directory is the same as the location of the UDDITaxonomyTools.jar then just the name is sufficient)

Filepath names should include the use of the forward-slash character (/) for all platforms.

For Cloudscape database users, the values of the following properties would be likely to be:

- classpath=c:/websphere/appserver/lib/db2j.jar; UDDITaxonomyTools.jar
- database.driver.className=com.ibm.db2j.jdbc.DB2jDriver
- database.url=jdbc:db2j:c:/ websphere/appserver/bin/uddi20

The string.delimiter is typically used where a description value contains the same character as the column delimiter character. For example, if the column.delimiter was set to ',' (comma), and there was a taxonomy

description value of 'Fruits, citrus', you could include this in the taxonomy data file by setting the string.delimiter to "(double quote) and enclosing the description in quotes: 'Fruits, citrus'. Note that the quote character is escaped with a backslash to indicate the literal character is to be used.

If a properties parameter is not specified, the utility looks for and uses configuration data set in a file called customTaxonomy.properties. This file is located, by default, in the <DeploymentManager-install-dir>/UDDIReg/scripts directory.

**Note:** to make updates to taxonomy data in a Cloudscape database, the IBM WebSphere Application Server must be stopped to release the connection to the database.

**Note:** There is currently a limitation with UDDITaxonomyTools.jar when used with a DB2 UDDI database and multi-byte characters such as Chinese, Japanese and Korean. The maximum number of multi-byte characters is the maximum value specified earlier for name, code, description and parentcode divided by 3. For example, name can only contain values up to 8 characters in length so the maximum number of Korean characters is 2. If the taxonomy file is found to have values that exceed the limits, a message is displayed by the tool indicating the line number and column where the problem occurs. This limitation does not affect use with a Cloudscape UDDI database.

### Publishing a Checked Categorization tModel

This section describes how to publish a checked categorization tModel with the 'customTaxonomy' keyedReferences to specify which custom taxonomy data to use and a display name.

**Note:** to specify an unchecked categorization substitute the 'checked' keyValue with 'unchecked' or, more simply, omit the keyedReference.

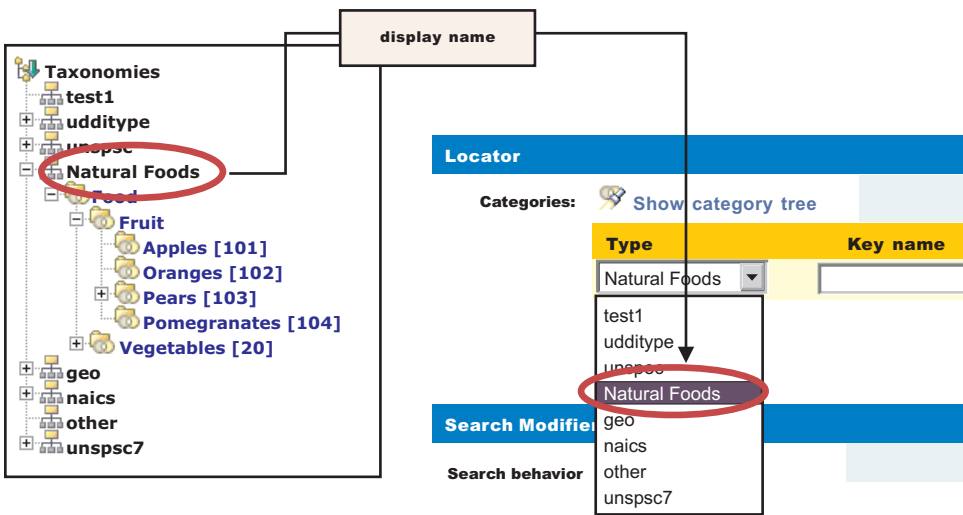
Publish a tModel to the IBM WebSphere UDDI Registry with a categoryBag containing keyedReferences as follows:

Note	tModelKey	KeyName	KeyValue
1	(uddi-org:types)	<optional>	categorization
2	(uddi-org:types)	<optional>	checked
3	(general keywords)	urn:x-ibm:uddi:customTaxonomy:key	<custom taxonomy name>
4	(general keywords)	urn:x-ibm:uddi:customTaxonomy:displayName	<custom taxonomy displayName>

1. Indicates this tModel is a categorization tModel (required)
2. Indicates use of the tModel will be checked against a list of valid data (required). (Omitting this keyedReference, or explicitly specifying a value of 'unchecked' will indicate this categorization is unchecked).
3. Indicates special use of the general keywords taxonomy, with a proprietary urn as the keyName value, defines the value used by the UDDI Registry to look up taxonomy data in its database. The value must be 1-8 (inclusive) characters long and corresponds directly with the name value in the custom taxonomy data file. Therefore, it must be unique within the registry.
4. Indicates special use of the general keywords taxonomy, with a proprietary urn as the keyName value, defines a name for the custom taxonomy that is intended for use in GUI implementations where the full tModel name might be too long\*. The value can be 1-255 characters (inclusive) long. If this keyedReference is not supplied, the name of the tModel should be used by the GUI implementation.

\* The displayName is intended to provide a way to label a taxonomy such that, when the UDDI GUI displays it in a taxonomy tree or in a pull-down list of available taxonomies, the meaning is clear to the user without being restricted to 8 characters and without needing to be the same as the published

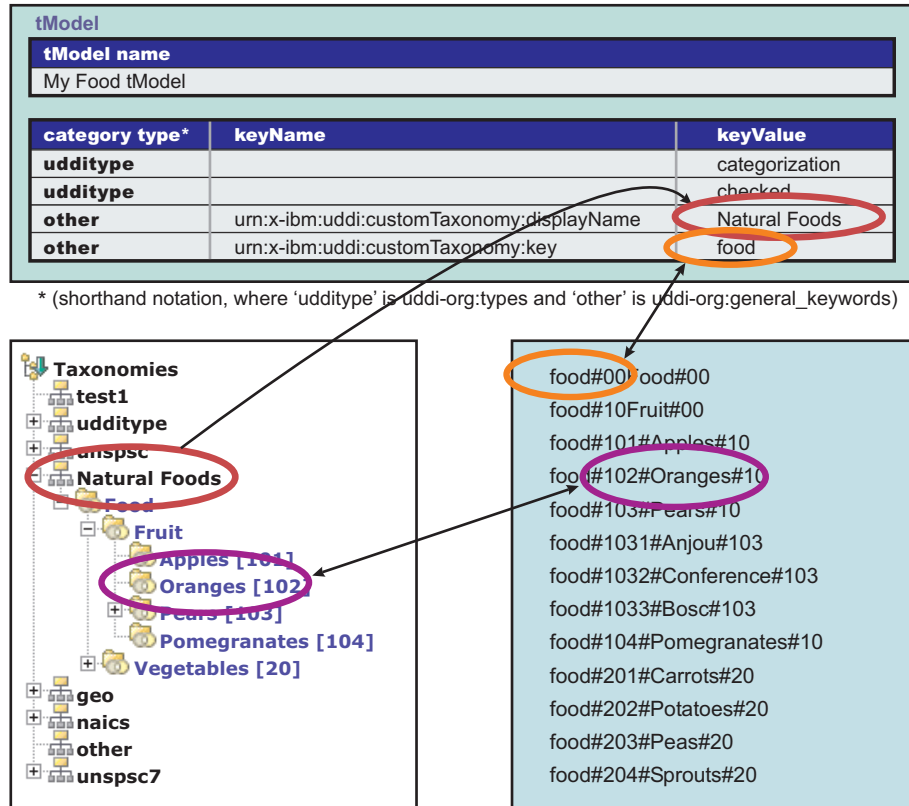
tModelName, which could be as long as 255 characters. An example is shown:



Uniqueness of the urn:x-ibm:uddi:customTaxonomy:key value is validated at the time a categorization tModel is published. If it is not unique, a UDDIInvalidValueException is returned. If using a GUI to publish the tModel, an appropriate message is displayed indicating the likely cause of the problem.

The urn:x-ibm:uddi:customTaxonomy:displayName should be unique if only to avoid confusion when displayed in GUIs but this is not validated.

The relationship between the various keyedReferences, the custom taxonomy data files and use in GUIs for a categorization tModel is shown below:



As a further example, to display the label 'Delicious Victuals' in GUI displays, the categorization tModel would have a keyedReference like this:

type	keyName	keyValue
other	urn:x-ibm:uddi:customTaxonomy:displayName	Delicious Victuals

And to link a categorization tModel to a custom taxonomy datafile with a taxonomy name of 'goodfood' the tModel's categoryBag must have a keyedReference like this:

type	keyName	keyValue
other	urn:x-ibm:uddi:customTaxonomy:key	goodfood

To publish a new categorization tModel using SOAP, the message would be:

```
<save_tModel generic="2.0" xmlns="urn:uddi-org:api_v2">
 <authInfo></authInfo>>
 <tModel tModelKey="">
 <name>Natural Foods tModel</name>
 <categoryBag>
 <keyedReference tModelKey="uuid:C1ACF26D-9672-4404-9D70-39B756E62AB4"
 keyValue="categorization"/>
 <keyedReference tModelKey="uuid:C1ACF26D-9672-4404-9D70-39B756E62AB4"
 keyValue="checked"/>
 <keyedReference tModelKey="uuid:A035A07C-F362-44DD-8F95-E2B134BF43B4"
 keyName="urn:x-ibm:uddi:customTaxonomy:key" keyValue="food"/>
 <keyedReference tModelKey="uuid:A035A07C-F362-44DD-8F95-E2B134BF43B4"
```

```

 keyName="urn:x-ibm:uddi:customTaxonomy:displayName" keyValue="Natural Foods"/>
 </categoryBag>
</tModel>
</save_tModel>

```

**Note:** 'uuid:C1ACF26D-9672-4404-9D70-39B756E62AB4' is the tModel key for uddi-org:types and 'uuid:A035A07C-F362-44DD-8F95-E2B134BF43B4' is the tModel key for uddi-org:general\_keywords.

## Validation and Error Handling

For a DB2-based IBM WebSphere UDDI Registry, custom taxonomy data can be loaded, removed and renamed using the provided utility without restarting the application (if you are using Cloudscape the application server must be stopped to make database updates). Removing data for which there is a corresponding checked categorization tModel will cause any use of that categorization's data to be reported as invalid.

**Note:** If an attempt is made to add data with a name that matches any of the 'internal' taxonomies, such as naics, geo, and so on, the request is rejected. If an attempt is made to rename or remove one of the internal taxonomies, a warning message is returned. Likewise if the user tries to rename a taxonomy to one of the reserved taxonomies, that is rejected.

The UDDI Registry user console performs validation while a save tModel request is being built, that is, before the publish occurs. For example, if a categorization tModel with a customTaxonomy:key keyValue of 'food' already exists (in a published categorization tModel), and the user tries to add a keyedReference with the same value to the current list of keyedReferences, the following message is displayed:

```
Advice: The 'urn:x-ibm:uddi:customTaxonomy:key' value of 'food' is already in use
by another categorization tModel. Enter a unique value
```

Similarly, only one of each of the customTaxonomy:key and customTaxonomy:displayName keyedReferences are allowed. For example, if the user tries to add two customTaxonomy:displayName keyedReferences the following message is displayed:

```
Advice: Only one 'urn:x-ibm:uddi:customTaxonomy:displayName' key name is allowed for the
'Other' taxonomy
```

If the customTaxonomy:key keyedReference is valid and unique at the time it is added to the save\_tModel request, the keyedReference is further validated when the user makes the publish request, to ensure that another session has not successfully published a categorization tModel with the same customTaxonomy:key. In this case, the user is returned to the Publish Technical Model page.

If a keyedReference containing a keyName value that starts with 'urn:x-ibm:uddi:customTaxonomy:' is followed by anything other than 'key' or 'displayName', the following message is displayed:

```
Advice: Only key name values of 'urn:x-ibm:uddi:customTaxonomy:displayName' and
'urn:x-ibm:uddi:customTaxonomy:key' are supported.
```

For SOAP, UDDI4J, and EJB initiated requests where the save\_tModel message may have multiple tModels, if any one of the tModels is a categorization tModel and it fails validation, the request fails with a UDDIInvalidValueException (plus additional information explaining the likely cause), and none of the tModels is published. For example, if a publish request includes a customTaxonomy:key keyedReference with a keyValue that matches the customTaxonomy:key keyValue of an existing categorization tModel, the following UDDIInvalidValueException is thrown, with the message:

```
E_invalidValue (20200) A value that was passed in a keyValue attribute did not pass validation.
This applies to checked categorizations, identifiers and other validated code lists. The error
text will clearly indicate the key and value combination that failed validation. Invalid
'customTaxonomy:dbKey' keyValue [naics] in keyedReference. KeyValue already in use by
tModelKey[UUID:C0B9FE13-179F-413D-8A5B-5004DB8E5BB2]
```



The customTaxonomy:key and customTaxonomy:displayName keyValue values are validated. For example, a publish categorization tModel request with a keyedReference including a customTaxonomy:key of 'toolongdbkey' was attempted, the following UDDIInvalidValueException is thrown, with the message:

```
E_invalidValue (20200) A value that was passed in a keyValue attribute did not pass validation.
This applies to checked categorizations, identifiers and other validated code lists. The error
text will clearly indicate the key and value combination that failed validation. Invalid
'customTaxonomy:key' keyValue [toolongdbkey] in keyedReference. tModelKey[]
```

If a categorization tModel is edited in the user console, or republished via SOAP, UDDI4J or EJB, such that it is no longer a categorization tModel (ie the categorization keyedReference is removed), then that tModel is removed from the internal store of categorization tModels, and its customTaxonomy:key value, if present, is available for use by new categorization tModels.

---

## UDDI Utility Tools

The UDDI Utility Tools is a suite of functions to help manage entities in a UDDI Registry. The functions allow users to:

- Search and select entities from a source UDDI Registry by specifying keys or search criteria
- Copy selected entities to a target UDDI Registry, while keeping the same entity key, and any child entities
- Delete selected entities from a target UDDI Registry
- Persist UDDI entities in an intermediate XML representation that can be used to customize and copy those entities to multiple target UDDI Registries
- Create new entities such as canonical tModels in a target UDDI Registry, including child entities
- Update existing entities in a target UDDI Registry, including child entities
- Automatically identify any entities referenced by an entity and import those into a target UDDI Registry first

UDDI Utility Tools can be used by running the UDDIUtilityTools.jar file, or from a batch file that invokes a specific Java class in the JAR file. UDDIUtilityTools.jar is located in the <DeploymentManager-install-dir>/UDDIReg/scripts directory. Alternatively, all of the functions of UDDI Utility Tools can be invoked through the supplied public Java API.

There are five main functions in UDDI Utility Tools:

### Export

Given an entity type and key, or a list of entity types and keys, UDDI Utility Tools gets the UDDI entities from the specified registry and writes them to the UDDI Entity Definition File. The entity type for each key can be one of business, service, bindingTemplate or tModel. The Entity Definition File contains XML that exactly describes each of the specified entities, according to the UDDI Utility Tools schema (which includes the UDDI V2 schema). The UDDI Entity Definition File separates entities by type, and automatically detects and records tModels referenced by the specified entities. You can use the 'referenced tModels' section of the file to ensure a target registry includes any referenced tModels before you try to import new entities to that registry.

### Import

Given a list of UDDI entities (which can be supplied using the UDDI Entity Definition File generated by the export function, possibly with additional editing, or programmatically in a container object), the import function detects if the entities already exist in the target registry and, if they do not, creates a minimal entity ("stub") with the specified key. The entities are then published updating the stubs with the supplied data and overwriting, or ignoring, existing entities as specified by the user. Note that the original key is maintained throughout.

### Promote

Combines the export and import steps such that the specified entities are extracted (by key) from the source registry and then imported into the target registry in a single logical step. The generation of a UDDI Entity Definition File is optional for this function.

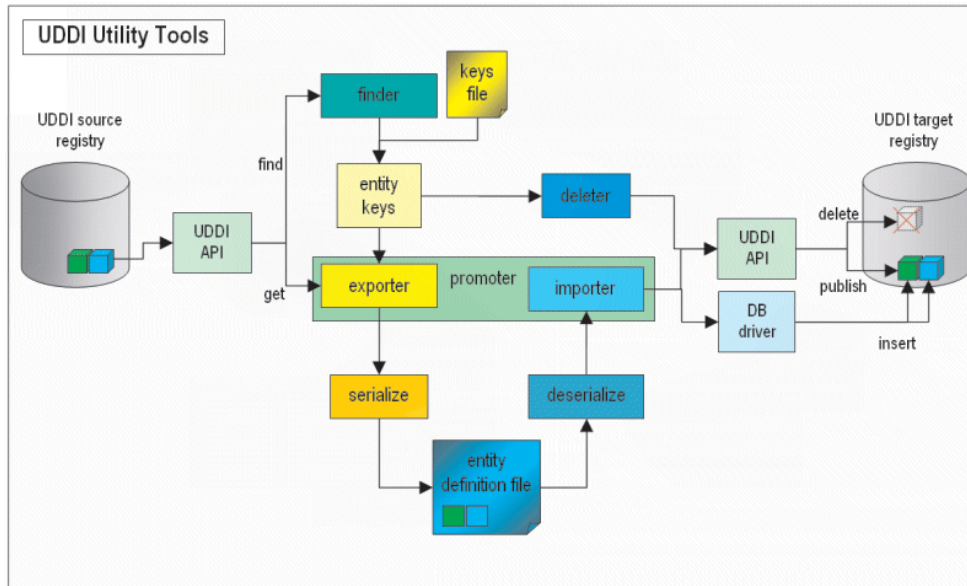
**Delete** Deletes the specified entities from the target UDDI Registry. The entities to delete are specified as an entity type, or a list of entity types, and keys, in the same way as for the export function.

**Find Matching Entities**

Takes as input search criteria in the form of UDDI Inquiry API objects for each of the various entity types. The set of entities that match the search criteria are used to generate a list of entity keys, and this in turn can be used as input to the export, promote and delete functions.

**Note:** This function is available only through the programmatic API.

The relationship between the functions, their input and output, and the source and target UDDI Registries is shown in this conceptual overview diagram:



**Setting up the configuration file**

Configuration data for UDDI Utility Tools resides in a configuration properties file, which describes the runtime environment, UDDI and database locations and access information, logging information, security configuration, entity definition file location, and other flags to control whether referenced entities are to be imported and/or overwritten.

UDDI Utility Tools is distributed with a sample configuration properties file (UDDIUtilityTools.properties) and this is searched for by default in the current directory if no properties path is specified. By default, this file is located in the <DeploymentManager-install-dir>/UDDIReg/scripts directory.

The most important property to set is the classpath, and this should include the current directory (.) and the UDDIUtilityTools.jar itself, plus all the dependent jars, most of which are located in the WebSphere AppServer lib directory. The classpath must include the database driver jar (for example db2java.zip). The other properties are well commented in the example properties file.

Below is an example properties file as distributed:

```
#####
Runtime environment
(if invoking via java -jar...)
"X Y" required around paths with spaces
#####
classpath=.;c:/Progra~1/WebSphere/utilitytools/UDDIUtilityTools.jar;
C:/WebSphere/AppServer/lib/soap.jar;C:/WebSphere/AppServer/java/jre/lib/ext/mail.jar;
C:/WebSphere/AppServer/java/jre/lib/ext/ibmjsse.jar;
```

```
C:/WebSphere/AppServer/java/jre/lib/ext/activation.jar;
C:/promoter/uddi4jv2.jar;C:/WebSphere/AppServer/lib/xerces.jar;
C:/WebSphere/AppServer/lib/j2ee.jar;"C:/Program Files/SQLLIB/java12/db2java.zip"
```

(ALL the above is on ONE line - shown here as broken for clarity)

```
#####
SOAP entry points for source UDDI
#####
fromInquiryURL=http://localhost:9080/uddisoap/inquiryapi
fromGetURL=http://localhost:9080/uddisoap/get

#####
SOAP entry points for target UDDI
#####
toInquiryURL=http://localhost:9080/uddisoap/inquiryapi
toPublishURL=http://localhost:9080/uddisoap/publishapi

#####
UDDI Registry user information
#
Note: this must match the user information
that was used to publish the entities on
the target UDDI registry.
#####
userID=UNAUTHENTICATED
password=NONE

#####
Configuration for destination UDDI DB
#####
dbDriver=COM.ibm.db2.jdbc.app.DB2Driver
dbUrl=jdbc:db2:uddi20
dbUser=db2admin (Your dbUser id)
dbPasswd=db2admin (Your dbPassword)

#####
Security provider configuration
#####
Indicates whether security is required on the target registry
secure.connection=true

The location of the truststore if security is required
trustStore.fileName=c:/websphere/appserver/etc/DummyClientTrustFile.jks

The password for the trust store
trustStore.password=WebAS

#####
Trace and message logging configuration
#####
detail level of message output (all functions)
verbose=true

detail level of trace output.
1: severe
2: normal
3: detail
traceLevel=3

path to message log file (relative or absolute)
messageLogFileName=logs/messages.log

path to trace log file (relative or absolute)
traceLogFileName=logs/trace.log

#####
```

```

Miscellaneous Options
#####
indicates if existing entities are overwritten (import/promote)
overwrite=false

indicates if referenced entities will be imported (import/promote)
importReferencedEntities=true

location of entity definition file, used for (export/import)
UddiEntityDefinitionFile=C:/definitions/entities01.xml

namespace prefix to use in definition file (export)
namespacePrefix=promote

```

## Prerequisites

UDDI Utility Tools must have the following jar files available. Their locations should be specified in the classpath property in the UDDI Utility Tools properties file:

### UDDIUtilityTools.jar

This is the tools jar itself. It MUST appear on the classpath.

### uddi4jv2.jar

This is the UDDI4J classes and can be found in <WAS\_HOME>/lib.

### db2java.zip

This is the DB2 JDBC driver classes jar which is provided with DB2.

### j2ee.jar

This contains some required J2EE classes.

### soap.jar

This is the Apache SOAP implementation.

### xerces.jar

This is the xerces XML parser.

Apache SOAP, and therefore UDDI Utility Tools, has a requirement of having activation.jar and mail.jar java extensions available. These should NOT be placed on the classpath but rather the ext folder of the JRE that is used to run the tool. If SSL is needed then ibmjssse.jar must also be in the ext folder. If you are using the JRE as supplied with IBM WebSphere Application Server then these extensions will already be in place and no further action is necessary.

The Security provider configuration section in the above properties file shows the location of the default DummyClientTrustFile.jks file. If you are using your own truststore, ensure that the location is placed here.

## The UDDI Entity Definition File

You generate this file by the export and promote functions, or you can choose to create it (either by hand, or by modifying a version of the file output by UDDI Utility Tools specifying the export function). It is the input to the import function.

**Note:** The extension to the uddi:tModel type to add a 'deleted' attribute is not currently used in UDDI Utility Tools.

The file is validated for well formedness and that it complies with the UDDI Utility Tools schema, shown here.

```

<?xml version="1.0" encoding="UTF-8" ?>
<xsd:schema id="uddiPromote" attributeFormDefault="unqualified" elementFormDefault="qualified"
 targetNamespace="http://www.ibm.com/xmlns/prod/WebSphere/UDDIUtilityTools"
 xmlns:xsd="http://www.w3.org/2001/XMLSchema"
 xmlns:uddi="urn:uddi-org:api_v2" xmlns="http://www.ibm.com/xmlns/prod/WebSphere/UDDIUtilityTools"

```

```

xmlns:promote="http://www.ibm.com/xmlns/prod/WebSphere/UDDIUtilityTools">

<xsd:import namespace="http://www.w3.org/XML/1998/namespace" schemaLocation="xml.xsd" />
<xsd:import namespace="urn:uddi-org:api_v2" schemaLocation="uddi_v2.xsd" />

<!-- define a type to represent state of a tModel -->
<xsd:simpleType name="tModelDeleted">
 <xsd:restriction base="xsd:NMTOKEN">
 <xsd:enumeration value="true" />
 <xsd:enumeration value="false" />
 </xsd:restriction>
</xsd:simpleType>

<!-- extend tModel with additional attribute of type tModelDeleted -->
<!-- This is restricted to values true or false -->
<xsd:complexType name="tModel">
 <xsd:complexContent>
 <xsd:extension base="uddi:tModel">
 <xsd:attribute name="deleted" type="promote:tModelDeleted" use="optional" />
 </xsd:extension>
 </xsd:complexContent>
</xsd:complexType>

<!-- Top level element definitions -->
<xsd:element name="uddiEntities" type="promote:uddiEntities" />
<xsd:complexType name="uddiEntities">
 <xsd:sequence>
 <xsd:element ref="promote:tModels" minOccurs="0" maxOccurs="1" />
 <xsd:element ref="promote:businesses" minOccurs="0" maxOccurs="1" />
 <xsd:element ref="promote:services" minOccurs="0" maxOccurs="1" />
 <xsd:element ref="promote:bindings" minOccurs="0" maxOccurs="1" />
 <xsd:element ref="promote:referencedTModels" minOccurs="0" maxOccurs="1" />
 </xsd:sequence>
</xsd:complexType>

<xsd:element name="businesses" type="promote:businesses" />
<xsd:complexType name="businesses">
 <xsd:sequence>
 <xsd:element ref="uddi:businessEntity" minOccurs="0" maxOccurs="unbounded" />
 </xsd:sequence>
</xsd:complexType>

<xsd:element name="tModels" type="promote:tModels" />
<xsd:complexType name="tModels">
 <xsd:sequence>
 <xsd:element ref="uddi:tModel" minOccurs="0" maxOccurs="unbounded" />
 </xsd:sequence>
</xsd:complexType>

<xsd:element name="services" type="promote:services" />
<xsd:complexType name="services">
 <xsd:sequence>
 <xsd:element ref="uddi:businessService" minOccurs="0" maxOccurs="unbounded" />
 </xsd:sequence>
</xsd:complexType>

<xsd:element name="bindings" type="promote:bindings" />
<xsd:complexType name="bindings">
 <xsd:sequence>
 <xsd:element ref="uddi:bindingTemplate" minOccurs="0" maxOccurs="unbounded" />
 </xsd:sequence>
</xsd:complexType>

<xsd:element name="referencedTModels" type="promote:referencedTModels" />
<xsd:complexType name="referencedTModels">
 <xsd:sequence>

```

```

 <xsd:element ref="uddi:tModel" minOccurs="0" maxOccurs="unbounded" />
 </xsd:sequence>
</xsd:complexType>
</xsd:schema>

```

## UDDI Entity Definition File example for canonical tModels

The example Entity Definition File following shows the five main sections for tModels, businesses, services, bindings and referencedTModels:

UDDI Utility Tools can be used to create new UDDI entities in a target UDDI Registry. A typical example of this is to introduce a new canonical tModel, which has a publicly known tModel key.

```

<?xml version="1.0" encoding="UTF-8"?>
<promote:uddiEntities xmlns="urn:uddi-org:api_v2"
 xmlns:promote="http://www.ibm.com/xmlns/prod/WebSphere/UDDIUtilityTools">

 <!-- tModels -->
 <promote:tModels>

 <tModel tModelKey="uuid:ee3966a8-faa5-416e-9772-128554343571" >
 <name>http://schemas.xmlsoap.org/ws/2002/07/policytmodel</name>
 <description>WS-PolicyAttachment policy expression</description>
 </tModel>

 <tModel tModelKey="uuid:ad61de98-4db8-31b2-a299-a2373dc97212" >
 <name>uddi-org:wSDL:address</name>
 <description xml:lang="en">
This tModel specifies the URL fact that the address obtains from the WSDL deployment file.
 </description>
 <overviewDoc>
 <overviewURL>
http://www.oasis-open.org/committees/uddi-spec/doc/tn/uddi-spec-tc-tn-wsdl-v2.htm#Address
 </overviewURL>
 </overviewDoc>
 </tModel>

 </promote:tModels>

 <!-- businesses -->
 <promote:businesses>
</promote:businesses>

 <!-- services -->
 <promote:services>
</promote:services>

 <!-- bindings -->
 <promote:bindings>
</promote:bindings>

 <!-- referenced tModels -->
 <promote:referencedTModels>
</promote:referencedTModels>

</promote:uddiEntities>

```

## Starting UDDI Utility Tools at a command prompt

Ensure the correct level of java is appropriate by setting the PATH statement to the level of java supplied with WebSphere. For example, from the command line, type:

On Windows:

```
set PATH=c:\Progra~1\WebSphere\AppServer\java\bin;%PATH%
```

On Unix and Linux platforms:

```
export PATH=/opt/WebSphere/AppServer/java/bin:$PATH
```

On Unix and Linux platforms run the db2profile script before issuing the java command to start UDDI Utility Tools. This script is located within the DB2 instance home directory under sqllib and is invoked by typing:

```
. /home/db2inst1/sqllib/db2profile
```

**Note:** In the above example, notice that the '.' is followed by a single space character.

**Note:** On Unix and Linux platforms the DB2 user **must** have a db2profile at \$HOME/sqllib/db2profile.

UDDI Utility Tools can be started using:

**java -jar UDDIUtilityTools.jar <function> [options]**

using a specified properties file that sets up classpath and other parameters, or it can be called using:

**java CommandLineProcessor**

where CommandLineProcessor is the class which processes command line arguments for UDDI Utility Tools, sets up configuration and invokes the appropriate function.

The usage is as follows:

```
Usage: java -jar UDDIUtilityTools.jar {function} [options]
```

function:

```
-promote <entity source> Promote entites between registries
-export <entity source> Extract entities from registry to XML
-delete <entity source> Delete entities from registry
-import Create entities from XML to registry
```

where <entity source> is one of:

```
-tmodel|-business|-service|-binding <key> Specify single entity type and key
-keysFile | -f <filename> Specify file containing entity types and keys
```

options:

```
-properties <filename> Specify path to configuration file
-overwrite | -o Overwrite an entity if it already exists
-log | -v Output verbose messages
-definitionFile <filename> Specify path to UDDI entity definition file
-importReferenced Import entities referenced by source entities
```

The following options override property settings in configuration file:

```
-overwrite
-log
-definitionFile
-importReferenced
```

```
Example: java -jar UDDIUtilityTools.jar -promote -keysFile C:/uddikeys.txt
```

Below are a set of UDDI Utility Tools command line examples:

To export a single business to the EDF file specified in a properties file in the current directory.

```
java -jar UDDIUtilityTools.jar -export -business 28B8B928-2B2E-4EC9-A647-1E40651E4752
```

As above but this time using a keys file to specify the entities to be exported

```
java -jar UDDIUtilityTools.jar -export -keysFile c:/myKeyFiles/keyFile01.txt
```

As above but also specifying verbose output to appear on the command line.

```
java -jar UDDIUtilityTools.jar -export -keysFile c:/myKeyFiles/keyFile02.txt -v
```



To import the contents of the default EDF specified in a UDDIUtilityTools.properties file in the current directory.

```
java -jar UDDIUtilityTools.jar -import
```

As above but also specifying that referenced tModels should be imported into the target registry.

```
java -jar UDDIUtilityTools.jar -import -importReferenced
```

To import the entities from an EDF at the specified location. Note the use of forward slashes even though this is an example on a windows file system.

```
java -jar UDDIUtilityTools.jar -import -definitionFile c:/myEDFs/entities01.xml
```

To import the entities from the default EDF including referenced tModels. Overwrite specifies that any entities excluding referenced tModels that are found in the target registry should be overwritten.

```
java -jar UDDIUtilityTools.jar -import -overwrite -importReferenced
```

To promote a single service from a source to a target registry using the properties file at a specified location.

```
java -jar UDDIUtilityTools.jar -promote -service 67961D67-330F-4F14-8210-E74A58E710F3
-properties c:/UUT/myUUTProps.properties
```

To promote a set of entities specified in a keys file.

```
java -jar UDDIUtilityTools.jar -promote -keysFile c:/myKeyFiles/keyFile03.txt
```

As above but specifying that existing entities in the target registry get overwritten.

```
java -jar UDDIUtilityTools.jar -promote -keysFile c:/myKeyFiles/keyFile04.txt -overwrite
```

To promote a set of entities specified in a keys file including referenced tModels.

```
java -jar UDDIUtilityTools.jar -promote -keysFile c:/myKeyFiles/keyFile05.txt -importReferenced
```

To promote a set of entities specified in a keys file but also create an EDF containing the promoted entities.

```
java -jar UDDIUtilityTools.jar -promote -keysFile c:/myKeyFiles/keyFile06.txt
-definitionFile c:/myEDFs/entities02.xml
```

To logically delete a single tModel. Note that it is not possible to physically delete tModels.

```
java -jar UDDIUtilityTools.jar -delete -tModel UUID:1E2B9D1E-E53D-4D36-9D46-6CCC176C466A
```

To delete all the entities specified in the keys file. Note that with the exception tModels all other entities will be physically deleted from the target registry.

```
java -jar UDDIUtilityTools.jar -delete -keysFile c:/myKeyFiles/keyFile04.txt
```

## A keys file example

Below is an example of the keys that are to be exported, promoted or deleted from the target registry:

```

Keys of entities to be exported, promoted from source registry or deleted from target registry

Note: keys must be comma separated and on SAME line
Note: property names are case sensitive. ('tmodels=' will be ignored)

businesses=97C77097-AC6C-4CA0-A6C4-452F7045C470, 4975E949-581F-4FCA-AD5F-E08280E05F9F
services=BB3864BB-1578-4833-8179-14391F14791F
bindings=
tModels=273F1727-7BFF-4FB5-A1FD-BA5C45BAFD9C
```

**Note:** If the importReferenced property is set to true, the list of tModels in the referencedTModels section is imported to the target registry. Minimal entities are created if the referencedTModel is new. If the referencedTModel already exists it is never overwritten, regardless of the overwrite property value. This is so that commonly referenced tModels such as categorization tModels do not keep being updated unnecessarily.

Should you need to update a referencedTModel, you must manually move the referencedTModel definition to the tModels section in the entity definition file and set overwrite to true.

## Content of the log files

Below shows examples of contents of two of the log files produced by running the tool. Note that some comments have been added in square brackets and in italic to highlight important points within the log file. The first is the messages.log which shows successful and unsuccessful operations for export, import and delete functions:

```
[18/09/03 22:57:24:684 BST] UDUT0002I: ***** Starting UDDI Utility Tools *****
[timestamp and eyecatcher indicate when tool is run]
[18/09/03 22:57:24:684 BST] UDUT0016I: Importing entities...
[18/09/03 22:57:24:684 BST] UDUT0031I: Deserializing...
[18/09/03 22:57:24:684 BST] UDUT0032I: Deserialized entities.
[18/09/03 22:57:24:684 BST] UDUT0124I: Created tModel minimal entity with tModelKey
[UUID:F2ACF26D-9672-4404-8E70-56B756E62CB2].
[18/09/03 22:57:24:684 BST] UDUT0115I: Imported 1 entities and 0 referenced entities.
[18/09/03 22:57:24:684 BST] UDUT0033I: Function '-import' completed successfully.
[this kind of message indicates the operation worked!]
[18/09/03 23:00:16:822 BST] UDUT0002I: ***** Starting UDDI Utility Tools *****
[18/09/03 23:00:16:822 BST] UDUT0009I: Exporting entities...
[18/09/03 23:00:16:822 BST] UDUT0015I: Exported 4 entities.
[18/09/03 23:00:16:822 BST] UDUT0016I: Importing entities...
[18/09/03 23:00:16:822 BST] UDUT0115I: Imported 1 entities and 0 referenced entities.
[18/09/03 23:00:16:822 BST] UDUT0033I: Function '-promote' completed successfully.
[18/09/03 23:06:17:741 BST] UDUT0002I: ***** Starting UDDI Utility Tools *****
[18/09/03 23:06:17:741 BST] UDUT0009I: Exporting entities...
[18/09/03 23:06:17:741 BST] UDUT0029I: Serializing...
[18/09/03 23:06:17:741 BST] UDUT0030I: Serialized entities.
[18/09/03 23:06:17:741 BST] UDUT0015I: Exported 1 entities.
[18/09/03 23:06:17:741 BST] UDUT0033I: Function '-export' completed successfully.
[18/09/03 23:07:15:624 BST] UDUT0002I: ***** Starting UDDI Utility Tools *****
[18/09/03 23:07:15:624 BST] UDUT0023I: Deleting entities...
[18/09/03 23:07:15:624 BST] UDUT0028I: Deleted 1 entities.
[18/09/03 23:07:15:624 BST] UDUT0033I: Function '-delete' completed successfully.
[18/09/03 23:09:43:326 BST] UDUT0002I: ***** Starting UDDI Utility Tools *****
[18/09/03 23:09:43:326 BST] UDUT0009I: Exporting entities...
[18/09/03 23:09:43:326 BST] UDUT0082E: Could not get service detail for
serviceKey[BB112CEE-3443-4617-BBCF-004C0A00BFFD].
```

The second log file shows a typical trace log file entry for an export:

```
[20/08/03 09:23:10:345 BST] ***** Starting UDDI Utility Tools *****
[eyecatcher and timestamp indicate when tool is run]
[20/08/03 09:23:10:345 BST] > com.ibm.uddi.promoter.config.ConfigurationValidator()
[the '>' indicates entry to the constructor of this class]
[20/08/03 09:23:10:345 BST] < com.ibm.uddi.promoter.config.ConfigurationValidator()
[the '<' indicates exit from the constructor]
[20/08/03 09:23:10:345 BST] > com.ibm.uddi.promoter.PromoterAPI()
[20/08/03 09:23:10:345 BST] com.ibm.uddi.promoter.PromoterAPI() Configuration:
[this configuration section reflects all the settings used by the tool]
TransformConfiguration:
 namespacePrefix=promote
 uddiEntityDefinitionFile=c:/definitions/testcycle.xml

ExportConfiguration:
 fromGetURL=http://localhost:9080/uddisoap/get
 fromInquiryURL=http://localhost:9080/uddisoap/inquiryapi
```

```
ImportConfiguration:
 overwrite=true
 uddiEntityDefinitionFile=c:/definitions/testcycle.xml
 importReferencedEntities=true
```

```
PublishConfiguration:
 toInquiryURL=http://localhost:9080/uddisoap/inquiryapi
 toPublishURL=http://localhost:9080/uddisoap/publishapi
 userID=UNAUTHENTICATED
 trustStoreFileName=XXX
 secureConnection=false
```

```
DatabaseConfiguration:
 dbDriver=COM.ibm.db2.jdbc.app.DB2Driver
 dbURL=jdbc:db2:uddi20
 dbUser=db2admin
```

```
LoggerConfiguration:
 messageStream=com.ibm.jvm.io.LocalizedPrintStream@2b233f6b
 messageLogFileName=c:/promoter/logs/messages.log
 traceLogFileName=c:/promoter/logs/trace.log
 traceLevel=4
 verbose=true
```

```
[20/08/03 09:23:10:345 BST] < com.ibm.uddi.promoter.PromoterAPI()
[20/08/03 09:23:10:345 BST] > com.ibm.uddi.promoter.PromoterAPI.setUddiEntities()
[20/08/03 09:23:10:345 BST] > com.ibm.uddi.promoter.export.KeyFileReader()
[20/08/03 09:23:10:345 BST] com.ibm.uddi.promoter.export.KeyFileReader() loaded tModel keys
 [log entries without a '>' or '<' are status messages only]
[20/08/03 09:23:10:345 BST] com.ibm.uddi.promoter.export.KeyFileReader() loaded business keys
[20/08/03 09:23:10:345 BST] com.ibm.uddi.promoter.export.KeyFileReader() loaded service keys
[20/08/03 09:23:10:345 BST] com.ibm.uddi.promoter.export.KeyFileReader() loaded binding keys
[20/08/03 09:23:10:345 BST] > com.ibm.uddi.promoter.UddiEntityKeys()
[20/08/03 09:23:10:345 BST] < com.ibm.uddi.promoter.UddiEntityKeys()
[20/08/03 09:23:10:345 BST] com.ibm.uddi.promoter.export.KeyFileReader() removed duplicate,
 empty and null keys
[20/08/03 09:23:10:345 BST] < com.ibm.uddi.promoter.export.KeyFileReader()
[20/08/03 09:23:10:345 BST] < com.ibm.uddi.promoter.PromoterAPI.setUddiEntities()
[20/08/03 09:23:10:345 BST] > com.ibm.uddi.promoter.PromoterAPI.exportEntities()
[20/08/03 09:23:10:345 BST] > com.ibm.uddi.promoter.export.EntityExtractor()
[20/08/03 09:23:10:345 BST] > com.ibm.uddi.promoter.TopologyUtils()
[20/08/03 09:23:10:345 BST] < com.ibm.uddi.promoter.TopologyUtils()
[20/08/03 09:23:10:345 BST] > com.ibm.uddi.promoter.export.ReferencedEntityExtractor()
[20/08/03 09:23:10:345 BST] < com.ibm.uddi.promoter.export.ReferencedEntityExtractor()
[20/08/03 09:23:10:345 BST] < com.ibm.uddi.promoter.export.EntityExtractor()
[20/08/03 09:23:10:345 BST] > com.ibm.uddi.promoter.UDDIClient()
[20/08/03 09:23:10:345 BST] com.ibm.uddi.promoter.UDDIClient() client type: 0
[20/08/03 09:23:10:345 BST] com.ibm.uddi.promoter.UDDIClient() userID: UNAUTHENTICATED
[20/08/03 09:23:10:345 BST] com.ibm.uddi.promoter.UDDIClient() inquiryURL:
 http://localhost:9080/uddisoap/inquiryapi
[20/08/03 09:23:10:345 BST] < com.ibm.uddi.promoter.UDDIClient()
[20/08/03 09:23:10:345 BST] > com.ibm.uddi.promoter.export.EntityExtractor.extractEntities()
[20/08/03 09:23:10:345 BST] > com.ibm.uddi.promoter.UddiEntities()
[20/08/03 09:23:10:345 BST] < com.ibm.uddi.promoter.UddiEntities()
[20/08/03 09:23:10:345 BST] > com.ibm.uddi.promoter.UDDIClient.getBusinessEntity()
[20/08/03 09:23:10:345 BST] < com.ibm.uddi.promoter.UDDIClient.getBusinessEntity()
[20/08/03 09:23:10:345 BST] > com.ibm.uddi.promoter.UDDIClient.getBusinessEntity()
[20/08/03 09:23:10:345 BST] < com.ibm.uddi.promoter.UDDIClient.getBusinessEntity()
[20/08/03 09:23:10:345 BST] > com.ibm.uddi.promoter.export.ReferencedEntityExtractor.discoverReferencedTModels()
[20/08/03 09:23:10:345 BST] > com.ibm.uddi.promoter.TopologyUtils.extractKeyedReferences()
[20/08/03 09:23:10:345 BST] < com.ibm.uddi.promoter.TopologyUtils.extractKeyedReferences()
[20/08/03 09:23:10:345 BST] > com.ibm.uddi.promoter.TopologyUtils.extractKeyedReferences()
[20/08/03 09:23:10:345 BST] < com.ibm.uddi.promoter.TopologyUtils.extractKeyedReferences()
[20/08/03 09:23:10:345 BST] > com.ibm.uddi.promoter.TopologyUtils.extractTModelsFromContacts()
[20/08/03 09:23:10:345 BST] < com.ibm.uddi.promoter.TopologyUtils.extractTModelsFromContacts()
```

```

[20/08/03 09:23:10:345 BST] <
 com.ibm.uddi.promoter.export.ReferencedEntityExtractor.discoverReferencedTModels()
[20/08/03 09:23:10:345 BST] > com.ibm.uddi.promoter.TopologyUtils.sortList()
[20/08/03 09:23:10:345 BST] < com.ibm.uddi.promoter.TopologyUtils.sortList()
[20/08/03 09:23:10:345 BST]
 com.ibm.uddi.promoter.export.EntityExtractor.extractEntities() removed discovery URLs
[20/08/03 09:23:10:345 BST]
 com.ibm.uddi.promoter.export.EntityExtractor.extractEntities() removed operator and authNames
[20/08/03 09:23:10:345 BST]
 com.ibm.uddi.promoter.export.EntityExtractor.extractEntities()
 retrieved UDDI entities from source registry
[20/08/03 09:23:10:345 BST] < com.ibm.uddi.promoter.export.EntityExtractor.extractEntities()
[20/08/03 09:23:10:345 BST] > com.ibm.uddi.promoter.transform.UddiSerializer()
[20/08/03 09:23:10:345 BST] < com.ibm.uddi.promoter.transform.UddiSerializer()
[20/08/03 09:23:10:345 BST] > com.ibm.uddi.promoter.transform.UddiSerializer.serialize()
[20/08/03 09:23:10:345 BST]
 com.ibm.uddi.promoter.transform.UddiSerializer.createTModelsSection() created tModels section
[20/08/03 09:23:10:345 BST]
 com.ibm.uddi.promoter.transform.UddiSerializer.createBusinessesSection() created businesses section
[20/08/03 09:23:10:345 BST]
 com.ibm.uddi.promoter.transform.UddiSerializer.createServicesSection() created services section
[20/08/03 09:23:10:345 BST]
 com.ibm.uddi.promoter.transform.UddiSerializer.createBindingsSection() created bindings section
[20/08/03 09:23:10:345 BST]
 com.ibm.uddi.promoter.transform.UddiSerializer.createReferencedTModelsSection()
 created referencedTModels section
[20/08/03 09:23:10:345 BST] < com.ibm.uddi.promoter.transform.UddiSerializer.serialize()
[20/08/03 09:23:10:345 BST] > com.ibm.uddi.promoter.UddiEntities.getEntityCount()
[20/08/03 09:23:10:345 BST]
 com.ibm.uddi.promoter.UddiEntities.getEntityCount() total entity count:1
[20/08/03 09:23:10:345 BST] < com.ibm.uddi.promoter.UddiEntities.getEntityCount()
[20/08/03 09:23:10:345 BST] < com.ibm.uddi.promoter.PromoterAPI.exportEntities()

```

## Starting UDDI Utility Tools through the API

UDDI Utility Tools provides a public API to functions for exporting, importing, promoting, finding and deleting UDDI entities. All of these functions can be invoked by using the PromoterAPI class. Usage of this class to perform these functions is typically to:

1. Create a Configuration object and populate it from a Properties object or from a configuration properties file.
2. Create a PromoterAPI object passing the Configuration in the constructor.
3. For keys based functions (export, delete and promote), set the keys by supplying a UDDIEntityKeys object, the location of the keys file, or, for one entity, by specifying an entity type and a key value.
4. Invoke the corresponding method for the function required: exportEntities, promoteEntities(boolean), importEntities, deleteEntities or extractKeysFromInquiry(FindTModel, FindBusiness, FindService, FindBinding, FindRelatedBusinesses).

The PromoterAPI classes should be sufficient for most users, but for users who want more specific control, classes are available for invoking lower level functions directly. See the Javadocs for more information on the interfaces.

Below are nine sets of sample code for UDDI Utility Tools demonstrating usage of the API classes. These are:

- TestEntityExporter.java
- TestEntityImporter.java
- TestEntityPromoter.java
- TestEntityFinder.java
- TestEntityDeleter.java
- TestUddiSerializer.java
- TestUddiDeserializer.java
- TestStubManager.java

- TestCreateMinimalEntity.java

## Known limitations with UDDI Utility Tools and workarounds

There are some known limitations with UDDI Utility Tools and a workaround for each. See UDDI troubleshooting tips for more information.

### Cloudfscape Restriction

The 'export' function when referencing a source registry with a Cloudfscape database is supported. However, the 'import', 'promote' and 'delete' functions are not supported when referencing a target registry because of a limitation with the UDDI Registry when working with a Cloudfscape database.

## UDDI TestEntityExporter.java

### TestEntityExporter.java:

```
package com.ibm.uddi.promoter.sample;

/*
 * IBM WebSphere UDDI Registry Sample Program
 *
 * "This sample program may be freely used, executed, copied and modified by customer
 * (a) for its own instruction and study, (b) in order to develop applications which run with
 * IBM WebSphere products, either for customer's own internal use or for redistribution with
 * customer's own products."
 *
 * Product 5630-A36, (C) COPYRIGHT International Business Machines Corp., 2003
 * All Rights Reserved * Licensed Materials - Property of IBM
 *
 */

import java.util.ArrayList;

import com.ibm.uddi.promoter.PromoterAPI;
import com.ibm.uddi.promoter.PromoterConstants;
import com.ibm.uddi.promoter.PromoterLogger;
import com.ibm.uddi.promoter.UddiEntities;
import com.ibm.uddi.promoter.UddiEntityKeys;
import com.ibm.uddi.promoter.config.Configuration;
import com.ibm.uddi.promoter.exception.PromoterConfigurationException;
import com.ibm.uddi.promoter.exception.PromoterException;
import com.ibm.uddi.promoter.exception.PromoterIOException;
import com.ibm.uddi.promoter.exception.PromoterTransportException;
import com.ibm.uddi.promoter.exception.PromoterUDDI4JException;
import com.ibm.uddi.promoter.export.EntityExtractor;

/**
 * Sample class to show typical use of the EntityExtractor
 * to export selected entities from a source registry.
 *
 * Exception handling is minimal.
 *
 * @author IBM
 */
public class TestEntityExporter {

 /**
 * Start this class with first argument value of:
 *
 * > 'exporterDirect' - shows how to use EntityExtractor directly
 * to perform export
 * > 'exporterAPI' - shows how to use PromoterAPI to perform export function
 *
 */
}
```

```

* @param args
*/
public static void main(String[] args) {

 if (args.length != 1
 || (args.length == 1 && !"exporterDirect".equals(args[0]) && !"exporterAPI"
 .equals(args[0]))) {
 System.out.println("must use either 'exporterDirect' or 'exporterAPI' as arg");
 } else if ("exporterDirect".equals(args[0])) {

 testExporterDirect();

 } else if ("exporterAPI".equals(args[0])) {

 testexporterViaPromoterAPI();

 }
}

/**
 * Instantiates PromoterAPI and invokes the exportEntities
 * method.
 *
 * The entity keys to be exported can be set up in one of three ways:
 *
 * 1. read the keys from a keys file.
 * 2. create and populate an UddiEntityKeys object
 * 3. specify an entity type and an entity key
 *
 * All three methods are shown.
 */
public static void testexporterViaPromoterAPI() {

 System.out.println("test export via PromoterAPI");

 // choose method of setting keys
 // (uncomment the other options to try them)
 //String keySetMethod = "keys file";

 //String keySetMethod = "single entity key";
 String keySetMethod = "keys object";

 try {

 // use a predefined config file - change the filepath for your setup
 Configuration config = new Configuration("c:/promoter/UDDIUtilityTools.properties");

 // write messages to console
 config.setMessageStream(System.out);

 // start trace and message logging
 PromoterLogger.getLogger().initialise(config.getLoggerConfiguration());

 // create a PromoterAPI to contain the keys
 PromoterAPI api = new PromoterAPI(config);

 // read the keys from a file
 if (keySetMethod.equals("keys file")) {

 api.setUddiEntities("c:/promoter/exportkeys.txt");

 // set up keys in UddiEntityKeys object
 } else if (keySetMethod.equals("keys object")) {

 // set up key values
 String businessKey1 = "1F5F411F-D33B-4BF1-B5AB-5FC7555FABD2";
 String serviceKey1 = "8998DA89-F9E3-438F-AB1E-2AB5DF2A1E32";

```

```

 // create lists for each type of entity
 ArrayList businessKeys = new ArrayList();
 ArrayList serviceKeys = new ArrayList();

 // add the keys to the lists for each entity type
 businessKeys.add(businessKey1);
 serviceKeys.add(serviceKey1);

 // create a container for all types of keys
 UddiEntityKeys entityKeys = new UddiEntityKeys();

 // add the business and service keys
 entityKeys.setBusinesses(businessKeys);
 entityKeys.setServices(serviceKeys);

 // set the keys in the PromoterAPI object
 api.setUddiEntities(entityKeys);

 // set up a single key for a specified entity type
 } else if (keySetMethod.equals("single entity key")) {

 // entity type - business, service, binding or tModel
 String entityType = PromoterConstants.ENTITY_BUSINESS;

 // key value
 String entityKey = "1F5F411F-D33B-4BF1-B5AB-5FC7555FABD2";

 api.setUddiEntity(entityType, entityKey);
 }

 // now perform export using the keys set
 api.exportEntities();

 System.out.println("finished exporting");
} catch (PromoterConfigurationException e) {
 System.out.println(e);
} catch (PromoterIOException e) {
 System.out.println(e);
} catch (PromoterException e) {
 System.out.println(e);
}
}

/**
 * Instantiates EntityExtractor and performs export for specified keys.
 *
 * Shows how to set up UddiEntityKeys, configure the EntityExtractor,
 * and invoke the exportEntities method.
 *
 * Note that this does not write the exported entities to a definition file - this example
 * shows how to get the UddiEntities. This can be used as input to UddiSerializer.
 */
public static void testExporterDirect() {

 System.out.println("test export using EntityExtractor");

 try {

 // set up key values
 String businessKey1 = "BEB2B1BE-F0E6-4619-95F6-00616400F66D";
 String serviceKey1 = "8998DA89-F9E3-438F-AB1E-2AB5DF2A1E32";

```



```

// create lists for each type of entity
ArrayList businessKeys = new ArrayList();
ArrayList serviceKeys = new ArrayList();

// add the keys to the lists for each type
businessKeys.add(businessKey1);
serviceKeys.add(serviceKey1);

// create a container for all types of keys
UddiEntityKeys entityKeys = new UddiEntityKeys();

// add the business and service keys
entityKeys.setBusinesses(businessKeys);
entityKeys.setServices(serviceKeys);

// use a predefined config file
Configuration config = new Configuration("c:/promoter/UDDIUtilityTools.properties");

// write messages to console
config.setOutputStream(System.out);

// start trace and message logging
PromoterLogger.getLogger().initialise(config.getLoggerConfiguration());

// create Entityexporter
EntityExtractor exporter = new EntityExtractor();

// set up export configuration
exporter.setExportConfig(config);

// set up the keys to be extracted from source registry
exporter.setEntityKeys(entityKeys);

// perform the export
UddiEntities extractedEntities = exporter.extractEntities();

System.out.println(extractedEntities.getEntityCount() + " entities were exported.");

System.out.println("finished exporting");

} catch (PromoterConfigurationException e) {
 System.out.println(e);

} catch (PromoterUDDI4JException e) {
 System.out.println(e);

} catch (PromoterTransportException e) {
 System.out.println(e);

} catch (PromoterException e) {
 System.out.println(e);
}
}
}

```

## UDDI TestEntityImporter.java

### TestEntityImporter.java

```

package com.ibm.uddi.promoter.sample;

/*
 * IBM WebSphere UDDI Registry Sample Program
 *
 * "This sample program may be freely used, executed, copied and modified by customer
 * (a) for its own instruction and study, (b) in order to develop applications which run with

```

```

* IBM WebSphere products, either for customer's own internal use or for redistribution with
* customer's own products."
*
* Product 5630-A36, (C) COPYRIGHT International Business Machines Corp., 2003
* All Rights Reserved * Licensed Materials - Property of IBM
*
*/

import java.util.ArrayList;
import java.util.Vector;

import org.uddi4j.datatype.Description;
import org.uddi4j.datatype.Name;
import org.uddi4j.datatype.tmodel.TModel;
import org.uddi4j.util.CategoryBag;
import org.uddi4j.util.KeyedReference;

import com.ibm.uddi.promoter.PromoterAPI;
import com.ibm.uddi.promoter.PromoterLogger;
import com.ibm.uddi.promoter.UddiEntities;
import com.ibm.uddi.promoter.config.Configuration;
import com.ibm.uddi.promoter.exception.PromoterConfigurationException;
import com.ibm.uddi.promoter.exception.PromoterException;
import com.ibm.uddi.promoter.exception.PromoterIOException;
import com.ibm.uddi.promoter.exception.PromoterTransportException;
import com.ibm.uddi.promoter.exception.PromoterUDDI4JException;
import com.ibm.uddi.promoter.publish.EntityImporter;

/**
 * Sample to show typical use of the EntityImporter
 * to import entity definitions to a target registry.
 *
 * @author IBM
 */
public class TestEntityImporter {

 /**
 * Start this class with first argument value of:
 *
 * > 'importerDirect' - shows how to use EntityImporter directly
 * to perform import
 * > 'importerAPI' - shows how to use PromoterAPI to perform import function
 *
 * @param args
 */

 public static void main(String[] args) {

 if (args.length != 1
 || (args.length == 1 && !"importerDirect".equals(args[0]) && !"importerAPI"
 .equals(args[0]))) {
 System.out.println("must use either 'importerDirect' or 'importerAPI' as arg");
 } else if ("importerDirect".equals(args[0])) {

 testImporterDirect();
 } else if ("importerAPI".equals(args[0])) {
 testImporterViaPromoterAPI();
 }
 }

 /**
 * Instantiates PromoterAPI and invokes the importEntities
 * method.
 *
 * The entities to be imported can be set up in one of two ways:
 *
 * 1. read the keys from a keys file.

```

```

* 2. create and populate an UddiEntityKeys object
* 3. specify an entity type and an entity key
*
*
* All three methods are shown.
*/
public static void testImporterViaPromoterAPI() {

 System.out.println("Starting testImportViaPromoterAPI");

 try {
 // use a predefined config file
 Configuration config = new Configuration("c:/promoter/UDDIUtilityTools.properties");

 // write messages to console
 config.setOutputStream(System.out);

 // start logging
 PromoterLogger.getLogger().initialise(config.getLoggerConfiguration());

 // override entity definition file location
 //config.setUddiEntityDefinitionFile("c:/definitions/entities01.xml");

 config.setOverwrite(true);

 // create a PromoterAPI to contain the keys
 PromoterAPI api = new PromoterAPI(config);

 // now perform import using the entity definition file
 api.importEntities();

 System.out.println("finished importing using PromoterAPI");

 } catch (PromoterConfigurationException e) {
 System.out.println(e);
 } catch (PromoterIOException e) {
 System.out.println(e);
 } catch (PromoterException e) {
 System.out.println(e);
 }
}

/**
 * Instantiates EntityImporter and performs import for entities in entity definition file.
 *
 * Shows how to set up the EntityImporter,
 * specify entities to be imported as UDDI4J objects (instead of using an EDF)
 * and invoke the importEntities method.
 */
public static void testImporterDirect() {

 System.out.println("Starting import using EntityImporter");

 try {

 // use a predefined config file
 Configuration config = new Configuration("c:/promoter/UDDIUtilityTools.properties");

 // write messages to console
 config.setOutputStream(System.out);

 // start logging
 PromoterLogger.getLogger().initialise(config.getLoggerConfiguration());

```

```

// overwrite existing data
config.setOverwrite(true);

// create EntityImporter
EntityImporter importer = new EntityImporter(config);

UddiEntities entities = new UddiEntities();

// set the key
String tModelKey = "UUID:1F5F411F-D33B-4BF1-B5AB-5FC7666FABA9";

// give the tModel a name
Name name = new Name("test tModel 1", "en");

// description
Vector descriptions = new Vector();
Description description = new Description();
description.setText("Example categorization tModel");
description.setLang("en-US");
descriptions.add(description);

// categoryBag
CategoryBag categoryBag = new CategoryBag();
KeyedReference keyedRef1 = new KeyedReference();
KeyedReference keyedRef2 = new KeyedReference();

// use Uddi types tModel (UUID:C1ACF26D-9672-4404-9D70-39B756E62AB4)
// to categorize this tModel as a checked categorization tModel
keyedRef1.setTModelKey(TModel.TYPES_TMODEL_KEY);
keyedRef1.setKeyName("Categorization (taxonomy)");
keyedRef1.setKeyValue("categorization");
keyedRef2.setTModelKey(TModel.TYPES_TMODEL_KEY);
keyedRef2.setKeyName("Checked value set");
keyedRef2.setKeyValue("checked");

categoryBag.add(keyedRef1);
categoryBag.add(keyedRef2);

// build the tModel to import
TModel tModel1 = new TModel();
tModel1.setName(name);
tModel1.setDescriptionVector(descriptions);
tModel1.setCategoryBag(categoryBag);
tModel1.setTModelKey(tModelKey);

ArrayList tModels = new ArrayList();
tModels.add(tModel1);

entities.setTModels(tModels);

importer.setEntities(entities);

// perform the import, and validate/sort referenced tModels
importer.importEntities(true);

System.out.println("finished importing using EntityImporter");
} catch (PromoterConfigurationException e) {
 System.out.println("pce: " + e);
} catch (PromoterUDDI4JException e) {
 System.out.println("ue: " + e);
} catch (PromoterTransportException e) {
 System.out.println("te: " + e);
} catch (PromoterException e) {
 System.out.println("pe: " + e);
}

```

```

 }
}
}

```

## UDDI TestEntityPromoter.java

### TestEntityPromoter.java:

```

package com.ibm.uddi.promoter.sample;

import java.util.ArrayList;

import com.ibm.uddi.promoter.PromoterAPI;
import com.ibm.uddi.promoter.PromoterConstants;
import com.ibm.uddi.promoter.PromoterLogger;
import com.ibm.uddi.promoter.UddiEntityKeys;
import com.ibm.uddi.promoter.config.Configuration;
import com.ibm.uddi.promoter.exception.PromoterConfigurationException;
import com.ibm.uddi.promoter.exception.PromoterException;
import com.ibm.uddi.promoter.exception.PromoterIOException;

/**
 * Sample to show typical use of the PromoterAPI
 * to export UDDI entities from a source registry and
 * import entity definitions to a target registry in one logical step.
 *
 * @author IBM
 */
public class TestEntityPromoter {

 public static void main(String[] args) {

 testPromote();
 }

 /**
 * Instantiates PromoterAPI and invokes the promoteEntities
 * method.
 *
 * The entities to be imported can be set up in one of two ways:
 *
 * 1. read the keys from a keys file.
 * 2. create and populate an UddiEntityKeys object
 * 3. specify an entity type and an entity key
 *
 * All three methods are shown.
 */
 public static void testPromote() {

 System.out.println("starting testPromote");

 // used to choose method of setting keys
 //String keySetMethod = "keys file";
 String keySetMethod = "keys object";
 //String keySetMethod = "single entity key";

 try {
 // use a predefined config file
 Configuration config = new Configuration("c:/promoter/UDDIUtilityTools.properties");

 // write messages to console
 config.setMessageStream(System.out);

 // start trace and message logging
 PromoterLogger.getLogger().initialise(config.getLoggerConfiguration());

```

```

// override entity definition file location
config.setUddiEntityDefinitionFile("c:/definitions/importTest.xml");

// create a PromoterAPI to contain the keys
PromoterAPI api = new PromoterAPI(config);

// read the keys from a file
if (keySetMethod.equals("keys file")) {

 api.setUddiEntities("c:/promoter/exportkeys.txt");

 // set up keys in UddiEntityKeys object
} else if (keySetMethod.equals("keys object")) {

 // set up key values
 String businessKey1 = "1F5F411F-D33B-4BF1-B5AB-5FC7555FABD2";
 String serviceKey1 = "3F5F411F-AAAA-4BF1-B5AB-5FC7555FABA3";

 // create lists for each type of entity
 ArrayList businessKeys = new ArrayList();
 ArrayList serviceKeys = new ArrayList();

 // add the keys to the lists for each type
 businessKeys.add(businessKey1);
 serviceKeys.add(serviceKey1);

 // create a container for all types of keys
 UddiEntityKeys entityKeys = new UddiEntityKeys();

 // add the business and service keys
 entityKeys.setBusinesses(businessKeys);
 entityKeys.setServices(serviceKeys);

 // set the keys in the PromoterAPI object
 api.setUddiEntities(entityKeys);

 // set up a single key for a specified entity type
} else if (keySetMethod.equals("single entity key")) {

 // entity type - business, service, binding or tModel
 String entityType = PromoterConstants.ENTITY_BUSINESS;

 // key value
 String entityKey = "1F5F411F-D33B-4BF1-B5AB-5FC7555FABD2";

 api.setUddiEntity(entityType, entityKey);
}

// now perform promote and save an EDF for reference (optional)
api.promoteEntities(true);

System.out.println("finished promoting");

} catch (PromoterConfigurationException e) {
 System.out.println(e);
} catch (PromoterIOException e) {
 System.out.println(e);
} catch (PromoterException e) {
 System.out.println(e);
}
}
}

```

# UDDI TestEntityFinder.java

## TestEntityFinder.java

```
:
package com.ibm.uddi.promoter.sample;

/*
 * IBM WebSphere UDDI Registry Sample Program
 *
 * "This sample program may be freely used, executed, copied and modified by customer
 * (a) for its own instruction and study, (b) in order to develop applications which run with
 * IBM WebSphere products, either for customer's own internal use or for redistribution with
 * customer's own products."
 *
 * Product 5630-A36, (C) COPYRIGHT International Business Machines Corp., 2003
 * All Rights Reserved * Licensed Materials - Property of IBM
 *
 */

import java.io.FileOutputStream;
import java.io.IOException;
import java.util.Properties;
import java.util.Vector;

import org.uddi4j.datatype.Name;
import org.uddi4j.request.FindBinding;
import org.uddi4j.request.FindBusiness;
import org.uddi4j.request.FindRelatedBusinesses;
import org.uddi4j.request.FindService;
import org.uddi4j.request.FindTModel;
import org.uddi4j.util.TModelBag;
import org.uddi4j.util.TModelKey;

import com.ibm.uddi.promoter.PromoterAPI;
import com.ibm.uddi.promoter.PromoterLogger;
import com.ibm.uddi.promoter.UddiEntityKeys;
import com.ibm.uddi.promoter.config.Configuration;
import com.ibm.uddi.promoter.exception.PromoterConfigurationException;
import com.ibm.uddi.promoter.exception.PromoterException;
import com.ibm.uddi.promoter.exception.PromoterIOException;
import com.ibm.uddi.promoter.exception.PromoterTransportException;
import com.ibm.uddi.promoter.exception.PromoterUDDI4JException;
import com.ibm.uddi.promoter.export.EntityFinder;

/**
 * Sample to show typical use of the EntityFinder to select keys and
 * generate keys.
 *
 * Be sure to set the MaxRows property of the FindXXXX objects with valid int values
 * as UDDI4J doesn't do any checking.
 *
 * @author IBM
 */
public class TestEntityFinder {

 /**
 * Start this class with first argument value of:
 *
 * > 'finderDirect' - shows how to use EntityFinder directly to perform find
 * > 'findererAPI' - shows how to use PromoterAPI to perform find function
 *
 * @param args
 */
 public static void main(String[] args) {
```



```

if (args.length != 1
 || (args.length == 1 && !"finderDirect".equals(args[0]) &&
 !"finderAPI".equals(args[0]))) {
 System.out.println("must use either 'finderDirect' or 'finderAPI' as arg");
} else if ("finderDirect".equals(args[0])) {

 testFinderDirect();
} else if ("finderAPI".equals(args[0])) {
 testFindViaPromoterAPI();
}
}

/**
 * Instantiates PromoterAPI and invokes the extractKeysFromInquiry
 * method using all the different find operations. It then invokes the
 * exportEntities function and writes an EDF, containing only those
 * entities that match the search criteria.
 */
public static void testFindViaPromoterAPI() {

 System.out.println("starting testFindViaPromoterAPI");

 try {
 // use a predefined config file
 Configuration config = new Configuration("c:/promoter/UDDIUtilityTools.properties");

 // write messages to console
 config.setMessageStream(System.out);

 // start trace and message logging
 PromoterLogger.getLogger().initialise(config.getLoggerConfiguration());

 PromoterAPI api = new PromoterAPI(config);

 // set up find criteria
 FindTModel findTModel = new FindTModel();
 findTModel.setName("u%");

 // this will force a truncation if there are more than 2 results
 findTModel.setMaxRows(2);

 FindBusiness findBusiness = new FindBusiness();
 Vector bizNames = new Vector();
 bizNames.add(new Name("%"));
 findBusiness.setNameVector(bizNames);

 // this means no truncation
 findBusiness.setMaxRows(0);

 FindService findService = new FindService();
 Vector serviceNames = new Vector();
 serviceNames.add(new Name("%"));
 findService.setNameVector(serviceNames);
 findService.setMaxRows(0);

 FindBinding findBinding = new FindBinding();
 TModelBag bag = new TModelBag();
 Vector tModelKeys = new Vector();
 // use your own key!
 findBinding.setServiceKey("1343C213-00FA-4A0F-8412-59FEC85912E8");

 tModelKeys.add(new TModelKey("UUID:8609C81E-EE1F-4D5A-B202-3EB13AD01823"));
 bag.setTModelKeyVector(tModelKeys);
 findBinding.setTModelBag(bag);
 findBinding.setMaxRows(0);
 }
}

```

```

FindRelatedBusinesses findRelatedBusinesses = new FindRelatedBusinesses();
findRelatedBusinesses.setBusinessKey("1F5F411F-AAAA-4BF1-B5AB-5FC7555FABA3");

// this means no truncation
findRelatedBusinesses.setMaxRows(0);

// this gets all the keys and invokes the setter in the API
api.extractKeysFromInquiry(findTModel, findBusiness, findService, null,
 findRelatedBusinesses);

// now perform export using the keys found
api.exportEntities();

 System.out.println("finished finding");
} catch (PromoterConfigurationException e) {
 System.out.println(e);
} catch (PromoterIOException e) {
 System.out.println(e);
} catch (PromoterException e) {
 System.out.println(e);
} catch (Exception e) {
 System.out.println(e);
}
}

/**
 * Instantiates EntityFinder and performs find for all types of entities.
 *
 * Checks truncated status and writes resultant keys list to the console.
 *
 * Also shows how to generate a keys file from UddiEntityKeys in format required
 * by export or delete.
 *
 * Typical output:
 * count: 6
 * tModel count: 2
 * combined business count: 1
 * service count: 2
 * binding count: 1
 * any results truncated? true
 * truncated tModels: true
 * truncated businesses: false
 * truncated services: false
 * truncated bindings: false
 * truncated related businesses: false
 * {services=3A90C43A-7745-457D-ADD6-1AB7071AD609, 1343C213-00FA-4A0F-8412-59FEC85912E8,
tModels=UUID:1A2B00BE-6E2C-42F5-875B-56F32686E0E7, UUID:5FCF5CD0-629A-4C50-8B16-F94E9CF2A674
bindings=BFC8DEBF-7BA4-444F-BB78-0518A8057828, businesses=8E0FA68E-C07A-4AF8-BE2A-A28164A22AAE}
 * finished
 *
 */
public static void testFinderDirect() {

 System.out.println("starting test of EntityFinder");

 try {

 // use a predefined config file
 Configuration config = new Configuration("c:/promoter/UDDIUtilityTools.properties");

 // write messages to console
 config.setMessageStream(System.out);

 // start trace and message logging
 PromoterLogger.getLogger().initialise(config.getLoggerConfiguration());

 // pass the config

```

```

EntityFinder finder = new EntityFinder(config);

// set up find criteria
FindTModel findTModel = new FindTModel();
findTModel.setName("u%");

// this will force a truncation if there are more than 2 results
findTModel.setMaxRows(2);

FindBusiness findBusiness = new FindBusiness();
Vector bizNames = new Vector();
bizNames.add(new Name("XXXXCCC"));
findBusiness.setNameVector(bizNames);

// this means no truncation
findBusiness.setMaxRows(0);

FindRelatedBusinesses findRelatedBusinesses = new FindRelatedBusinesses();
findRelatedBusinesses.setBusinessKey("330D4833-850A-4ADD-BE4A-A00264A04AF0");

// this means no truncation
findRelatedBusinesses.setMaxRows(0);

FindService findService = new FindService();
Vector serviceNames = new Vector();
serviceNames.add(new Name("%"));
findService.setNameVector(serviceNames);
findService.setMaxRows(0);

FindBinding findBinding = new FindBinding();
TModelBag bag = new TModelBag();
Vector tModelKeys = new Vector();
// use your own key!
findBinding.setServiceKey("1343C213-00FA-4A0F-8412-59FEC85912E8");

tModelKeys.add(new TModelKey("UUID:8609C81E-EE1F-4D5A-B202-3EB13AD01823"));
bag.setTModelKeyVector(tModelKeys);
findBinding.setTModelBag(bag);
findBinding.setMaxRows(0);

// perform the find and get the keys back
UddiEntityKeys keys =
 finder.findEntities(findTModel, findBusiness, findService, findBinding,
 findRelatedBusinesses);

// output what we got
System.out.println("count: " + keys.getKeyCount());
System.out.println("tModel count: " + keys.getTModels().size());
System.out.println("combined business count: " + keys.getBusinesses().size());
System.out.println("service count: " + keys.getServices().size());
System.out.println("binding count: " + keys.getBindings().size());

// Test the TruncatedStates methods
// use quick test first
System.out.println("any results truncated? " + finder.getTruncatedStates().
 isAnyStateTruncated());

// then look at each result type
System.out.println("truncated tModels: " + finder.getTruncatedStates().
 isTModelResultsTruncated());
System.out.println("truncated businesses: " + finder.getTruncatedStates().
 isBusinessResultsTruncated());
System.out.println("truncated services: " + finder.getTruncatedStates().
 isServiceResultsTruncated());
System.out.println("truncated bindings: " + finder.getTruncatedStates().
 isBindingResultsTruncated());
System.out.println(

```

```

 "truncated related businesses: " + finder.getTruncatedStates().
 isRelatedBusinessResultsTruncated());

 // what Properties does UddiEntityKeys return?
 Properties properties = keys.getKeysAsProperties();
 System.out.println(properties.toString());

 // write the keys to a properties file in format required for a UUT keys file
 FileOutputStream out = new FileOutputStream("c:/promoter/generatedKeysFile.txt");
 keys.writeKeysAsPropertiesFile(out, ">>> This is my custom header <<<");
 out.close();

 System.out.println("finished EntityFinder test");

} catch (PromoterConfigurationException e) {

 System.out.println("pce: " + e);
} catch (PromoterUDDI4JException e) {
 System.out.println("ue: " + e);
} catch (PromoterTransportException e) {
 System.out.println("te: " + e);
} catch (PromoterException e) {
 System.out.println("pe: " + e);
} catch (IOException e) {
 System.out.println("ioe: " + e);
}
}
}

```

## UDDI TestEntityDeleter.java

### TestEntityDeleter-java

```

package com.ibm.uddi.promoter.sample;

/*
 * IBM WebSphere UDDI Registry Sample Program
 *
 * "This sample program may be freely used, executed, copied and modified by customer
 * (a) for its own instruction and study, (b) in order to develop applications which run with
 * IBM WebSphere products, either for customer's own internal use or for redistribution with
 * customer's own products."
 *
 * Product 5630-A36, (C) COPYRIGHT International Business Machines Corp., 2003
 * All Rights Reserved * Licensed Materials - Property of IBM
 *
 */

import java.util.ArrayList;

import com.ibm.uddi.promoter.PromoterAPI;
import com.ibm.uddi.promoter.PromoterConstants;
import com.ibm.uddi.promoter.PromoterLogger;
import com.ibm.uddi.promoter.UddiEntityKeys;
import com.ibm.uddi.promoter.config.Configuration;
import com.ibm.uddi.promoter.exception.PromoterConfigurationException;
import com.ibm.uddi.promoter.exception.PromoterException;
import com.ibm.uddi.promoter.exception.PromoterIOException;
import com.ibm.uddi.promoter.exception.PromoterTransportException;
import com.ibm.uddi.promoter.exception.PromoterUDDI4JException;
import com.ibm.uddi.promoter.publish.EntityDeleter;

/**
 * Sample to show typical use of the EntityDeleter
 * to delete selected entities from a target registry.
 *
 * Exception handling is minimal.
 */

```

```

*
* @author IBM
*/
public class TestEntityDeleter {

 /**
 * Start this class with first argument value of:
 *
 * > 'deleterDirect' - shows how to use EntityDeleter directly to perform
 * delete
 * > 'deleterAPI' - shows how to use PromoterAPI to perform delete function
 *
 * @param args
 */
 public static void main(String[] args) {

 if (args.length != 1
 || (args.length == 1 && !"deleterDirect".equals(args[0]) &&
 !"deleterAPI".equals(args[0]))) {
 System.out.println("must use either 'deleterDirect' or 'deleterAPI' as arg");
 } else if ("deleterDirect".equals(args[0])) {

 testDeleterDirect();
 } else if ("deleterAPI".equals(args[0])) {
 testDeleterViaPromoterAPI();
 }
 }

 /**
 * Instantiates PromoterAPI and invokes the deleteEntities
 * method.
 *
 * The entity keys to be deleted can be set up in one of three ways:
 *
 * 1. read the keys from a keys file.
 * 2. create and populate an UddiEntityKeys object
 * 3. specify an entity type and an entity key
 *
 * All three methods are shown.
 */
 public static void testDeleterViaPromoterAPI() {

 System.out.println("starting testDeleteViaPromoterAPI");

 // used to choose method of setting keys
 //String keySetMethod = "keys file";

 //String keySetMethod = "keys object";
 String keySetMethod = "single entity key";

 try {
 // use a predefined config file
 Configuration config = new Configuration("c:/promoter/UDDIUtilityTools.properties");

 // write messages to console
 config.setOutputStream(System.out);

 // start trace and message logging
 PromoterLogger.getLogger().initialise(config.getLoggerConfiguration());

 // create a PromoterAPI to contain the keys
 PromoterAPI api = new PromoterAPI(config);

 // read the keys from a file
 if (keySetMethod.equals("keys file")) {

```

```

 api.setUddiEntities("c:/promoter/deletekeys.txt");

 // set up keys in UddiEntityKeys object
 } else if (keySetMethod.equals("keys object")) {

 // set up key values
 String businessKey1 = "BEB2B1BE-F0E6-4619-95F6-00616400F66D";
 String serviceKey1 = "1F5F411F-D33B-4BF1-B5AB-5FC7555FABD2";

 // create lists for each type of entity
 ArrayList businessKeys = new ArrayList();
 ArrayList serviceKeys = new ArrayList();

 // add the keys to the lists for each type
 businessKeys.add(businessKey1);
 serviceKeys.add(serviceKey1);

 // create a container for all types of keys
 UddiEntityKeys entityKeys = new UddiEntityKeys();

 // add the business and service keys
 entityKeys.setBusinesses(businessKeys);
 entityKeys.setServices(serviceKeys);

 // set the keys in the PromoterAPI object
 api.setUddiEntities(entityKeys);

 // set up a single key for a specified entity type
 } else if (keySetMethod.equals("single entity key")) {

 // entity type - business, service, binding or tModel
 String entityType = PromoterConstants.ENTITY_BUSINESS;

 // key value
 String entityKey = "1F5F411F-D33B-4BF1-B5AB-5FC7555FABA3";

 api.setUddiEntity(entityType, entityKey);
 }

 // now perform delete using the keys set
 api.deleteEntities();

 System.out.println("finished deleting using PromoterAPI");

} catch (PromoterConfigurationException e) {
 System.out.println(e);
} catch (PromoterIOException e) {
 System.out.println(e);
} catch (PromoterException e) {
 System.out.println(e);
}
}

/**
 * Instantiates EntityDeleter and performs delete for specified keys.
 *
 * Shows how to set up UddiEntityKeys, configure the EntityDeleter,
 * and invoke the deleteEntities method.
 */
public static void testDeleterDirect() {

 System.out.println("starting testDeleterDirect");

 try {

```

```

// set up key values
String businessKey1 = "BEB2B1BE-F0E6-4619-95F6-00616400F66D";
String serviceKey1 = "1F5F411F-D33B-4BF1-B5AB-5FC7555FABD2";

// create lists for each type of entity
ArrayList businessKeys = new ArrayList();
ArrayList serviceKeys = new ArrayList();

// add the keys to the lists for each type
businessKeys.add(businessKey1);
serviceKeys.add(serviceKey1);

// create a container for all types of keys
UddiEntityKeys entityKeys = new UddiEntityKeys();

// add the business and service keys
entityKeys.setBusinesses(businessKeys);
entityKeys.setServices(serviceKeys);

// use a predefined config file
Configuration config = new Configuration("c:/promoter/UDDIUtilityTools.properties");

// write messages to console
config.setMessageStream(System.out);

// start trace and message logging
PromoterLogger.getLogger().initialise(config.getLoggerConfiguration());

// create EntityDeleter
EntityDeleter deleter = new EntityDeleter(entityKeys);

// set up deleter configuration
deleter.setDeleteConfig(config);

// perform the delete
deleter.deleteEntities();

System.out.println("finished deleting using EntityDeleter");

} catch (PromoterConfigurationException e) {

 System.out.println("pce: " + e);
} catch (PromoterUDDI4JException e) {
 System.out.println("ue: " + e);
} catch (PromoterTransportException e) {
 System.out.println("te: " + e);
} catch (PromoterException e) {
 System.out.println("pe: " + e);
}
}
}

```

## UDDI TestUddiSerializer.java

### TestUddiSerializer.java:

```

package com.ibm.uddi.promoter.sample;

import java.util.ArrayList;
import java.util.Vector;

import org.uddi4j.datatype.Description;
import org.uddi4j.datatype.Name;
import org.uddi4j.datatype.tmodel.TModel;
import org.uddi4j.util.CategoryBag;

```



```

import org.uddi4j.util.KeyedReference;

import com.ibm.uddi.promoter.PromoterLogger;
import com.ibm.uddi.promoter.UddiEntities;
import com.ibm.uddi.promoter.config.Configuration;
import com.ibm.uddi.promoter.exception.PromoterConfigurationException;
import com.ibm.uddi.promoter.exception.PromoterException;
import com.ibm.uddi.promoter.exception.PromoterIOException;
import com.ibm.uddi.promoter.transform.UddiSerializer;

/**
 * Sample to show how the UddiSerializer can be used to
 * extract UDDI entities from a source registry and write their
 * definitions to an entity definition file.
 *
 * @author IBM
 */
public class TestUddiSerializer {

 public static void main(String[] args) {

 testSerializer();
 }

 /**
 * Instantiates UddiSerializer and invokes the serialize
 * method.
 */
 public static void testSerializer() {

 System.out.println("starting serialization");

 // used to choose method of setting keys
 String keySetMethod = "keys file";

 try {
 // use a predefined config file
 Configuration config = new Configuration("c:/promoter/UDDIUtilityTools.properties");

 // write messages to console
 config.setMessageStream(System.out);

 // start trace and message logging
 PromoterLogger.getLogger().initialise(config.getLoggerConfiguration());

 UddiSerializer serializer = new UddiSerializer(config.getTransformConfiguration());

 UddiEntities entities = new UddiEntities();

 // we set the key
 String tModelKey = "UUID:F2ACF46D-9672-4404-8E70-56B756E62CB2";

 // give the tModel a name
 Name name = new Name("test tModel 1", "en");

 // description
 Vector descriptions = new Vector();
 Description description = new Description();
 description.setText("Example categorization tModel");
 description.setLang("en-US");
 descriptions.add(description);

 // categoryBag
 CategoryBag categoryBag = new CategoryBag();
 KeyedReference keyedRef1 = new KeyedReference();
 KeyedReference keyedRef2 = new KeyedReference();

```

```

// use Uddi types tModel (UUID:C1ACF26D-9672-4404-9D70-39B756E62AB4)
// to categorize this tModel as a checked categorization tModel
keyedRef1.setTModelKey(TModel.TYPES_TMODEL_KEY);
keyedRef1.setKeyName("Categorization (taxonomy)");
keyedRef1.setKeyValue("categorization");
keyedRef2.setTModelKey(TModel.TYPES_TMODEL_KEY);
keyedRef2.setKeyName("Checked value set");
keyedRef2.setKeyValue("checked");

categoryBag.add(keyedRef1);
categoryBag.add(keyedRef2);

// build the tModel to save in the EDF
TModel tModel1 = new TModel();
tModel1.setName(name);
tModel1.setDescriptionVector(descriptions);
tModel1.setCategoryBag(categoryBag);
tModel1.setTModelKey(tModelKey);

ArrayList tModels = new ArrayList();
tModels.add(tModel1);

entities.setTModels(tModels);

// write the EDF
serializer.serialize(entities);

System.out.println("finished writing EDF");

} catch (PromoterConfigurationException e) {
 System.out.println(e);
} catch (PromoterIOException e) {
 System.out.println(e);
} catch (PromoterException e) {
 System.out.println(e);
}
}
}
}

```

## UDDI TestUddiDeserializer.java

### TestUddiDeserializer.java:

```

package com.ibm.uddi.promoter.sample;

import java.util.Iterator;
import java.util.List;
import java.util.Vector;

import org.uddi4j.datatype.Name;
import org.uddi4j.datatype.business.BusinessEntity;

import com.ibm.uddi.promoter.PromoterLogger;
import com.ibm.uddi.promoter.UddiEntities;
import com.ibm.uddi.promoter.config.Configuration;
import com.ibm.uddi.promoter.exception.PromoterConfigurationException;
import com.ibm.uddi.promoter.exception.PromoterException;
import com.ibm.uddi.promoter.exception.PromoterIOException;
import com.ibm.uddi.promoter.transform.UddiDeserializer;

/**
 * Sample to show how the UddiDeserializer can be used to
 * read from an entity definition file and create UDDI4J objects.
 *
 * @author IBM

```

```

*/
public class TestUddiDeserializer {

 public static void main(String[] args) {

 testDeserializer();
 }

 /**
 * Instantiates UddiDeserializer and invokes the deserialize
 * method.
 */
 public static void testDeserializer() {

 System.out.println("starting testDeserializer");

 try {
 // use a predefined config file
 Configuration config = new Configuration("c:/promoter/UDDIUtilityTools.properties");

 // write messages to console
 config.setMessageStream(System.out);

 // start trace and message logging
 PromoterLogger.getLogger().initialise(config.getLoggerConfiguration());

 // override the location of the EDF
 config.setUddiEntityDefinitionFile("c:/definitions/testDeserialize.xml");

 UddiDeserializer deserializer = new UddiDeserializer(config.getTransformConfiguration());

 // read the EDF
 UddiEntities entities = deserializer.deserialize();

 List businesses = entities.getBusinesses();

 // just display the businesses, for this example
 for (Iterator iter = businesses.iterator(); iter.hasNext();) {

 BusinessEntity business = (BusinessEntity) iter.next();

 System.out.println("business key: " + business.getBusinessKey());

 Vector names = business.getNameVector();

 for (Iterator namesIterator = names.iterator(); namesIterator.hasNext();) {
 Name name = (Name) namesIterator.next();
 System.out.println("business name: " + name.getText());
 }
 }

 System.out.println("finished reading EDF");

 } catch (PromoterConfigurationException e) {
 System.out.println(e);
 } catch (PromoterIOException e) {
 System.out.println(e);
 } catch (PromoterException e) {
 System.out.println(e);
 }
 }
}

```

## UDDI TestStubManager.java

### TestStubManager.java:

```
package com.ibm.uddi.promoter.sample;

/*
 * IBM WebSphere UDDI Registry Sample Program
 *
 * "This sample program may be freely used, executed, copied and modified by customer
 * (a) for its own instruction and study, (b) in order to develop applications which run with
 * IBM WebSphere products, either for customer's own internal use or for redistribution with
 * customer's own products."
 *
 * Product 5630-A36, (C) COPYRIGHT International Business Machines Corp., 2003
 * All Rights Reserved * Licensed Materials - Property of IBM
 *
 */

import java.sql.Connection;
import java.util.Vector;

import org.uddi4j.datatype.Name;
import org.uddi4j.datatype.business.BusinessEntity;

import com.ibm.uddi.promoter.PromoterLogger;
import com.ibm.uddi.promoter.Publisher;
import com.ibm.uddi.promoter.config.Configuration;
import com.ibm.uddi.promoter.config.DatabaseConfiguration;
import com.ibm.uddi.promoter.db.DBManager;
import com.ibm.uddi.promoter.exception.PromoterDBException;
import com.ibm.uddi.promoter.exception.PromoterException;
import com.ibm.uddi.promoter.publish.EntityStubManager;
import com.ibm.uddi.promoter.publish.KeyDetector;

/**
 * @author IBM
 *
 * Example of how to write minimal entities to UDDI database using StubManager
 */
public class TestStubManager {

 public static void main(String[] args) {

 System.out.println("Inserting minimal entity using StubManager");

 try {
 Configuration config = null;

 // get config data from properties file
 // (alternatively this object can be set programmatically with a Properties object)
 config = new Configuration("c:/promoter/UDDIUtilityTools.properties");

 // write messages to console
 config.setMessageStream(System.out);

 // start trace and message logging
 PromoterLogger.getLogger().initialise(config.getLoggerConfiguration());

 DatabaseConfiguration dbConfig = config.getDatabaseConfiguration();

 // get database manager
 DBManager dbManager = DBManager.getInstance(dbConfig);

 // and a connection
 Connection connection = dbManager.getConnection();
```

```

// KeyDetector ensures no duplicate keys get into database
KeyDetector keyDetector = new KeyDetector(connection);

Publisher publisher = new Publisher();
publisher.setOperatorName("test operator name");
publisher.setAuthorizedName("UNAUTHENTICATED");

EntityStubManager stubManager = new EntityStubManager(publisher, keyDetector);

// build a minimal businessEntity
BusinessEntity business = new BusinessEntity();
business.setBusinessKey("1F6F411F-D33B-4BF1-B5AB-5FC7555FABA2");
Vector names = new Vector();
Name name = new Name("test business 1", "en-US");

names.add(name);
business.setNameVector(names);

// insert the business
stubManager.insertBusiness(business);

System.out.println("finished inserting minimal entity");

// this is likely to be thrown if the entity already exists
} catch (PromoterDBException e) {
 System.out.println(e);
} catch (PromoterException e) {
 System.out.println(e);
}
}
}
}

```

## UDDI TestCreateMinimalEntity.java

### TestCreateMinimalEntity.java:

```

package com.ibm.uddi.promoter.sample;

/*
 * IBM WebSphere UDDI Registry Sample Program
 *
 * "This sample program may be freely used, executed, copied and modified by customer
 * (a) for its own instruction and study, (b) in order to develop applications which run with
 * IBM WebSphere products, either for customer's own internal use or for redistribution with
 * customer's own products."
 *
 * Product 5630-A36, (C) COPYRIGHT International Business Machines Corp., 2003
 * All Rights Reserved * Licensed Materials - Property of IBM
 */

import java.sql.Connection;

import com.ibm.uddi.promoter.PromoterLogger;
import com.ibm.uddi.promoter.config.Configuration;
import com.ibm.uddi.promoter.config.DatabaseConfiguration;
import com.ibm.uddi.promoter.db.DBException;
import com.ibm.uddi.promoter.db.DBManager;
import com.ibm.uddi.promoter.entity.BusinessStub;
import com.ibm.uddi.promoter.entity.ServiceStub;
import com.ibm.uddi.promoter.exception.PromoterDBException;
import com.ibm.uddi.promoter.exception.PromoterException;
import com.ibm.uddi.promoter.publish.KeyDetector;

/**
 * Example of how to write minimal data entities directly to UDDI database.

```

```

*
* It is up to the developer to update the entity using UDDI4J publish operations with
* further information. Normally there would be no need to use these classes, as
* the PromoterAPI will handle minimal entity creation and updates.
*
* @author IBM
*/
public class TestCreateMinimalEntity {

 public static void main(String[] args) {

 System.out.println("Writing minimal entities to UDDI registry using BusinessStub and ServiceStub classes");

 try {
 Configuration config = null;
 // get config data from properties file (could easily be set programmatically)
 config = new Configuration("c:/promoter/UDDIUtilityTools.properties");

 // write messages to console
 config.setOutputStream(System.out);

 // start trace and message logging
 PromoterLogger.getLogger().initialise(config.getLoggerConfiguration());

 DatabaseConfiguration dbConfig = config.getDatabaseConfiguration();

 // get db manager
 DBManager dbManager = DBManager.getInstance(dbConfig);

 Connection connection = dbManager.getConnection();

 // KeyDetector ensures no duplicate keys get into database
 KeyDetector keyDetector = new KeyDetector(connection);

 // create minimal entity object (business)
 BusinessStub businessSeed = new BusinessStub(keyDetector);

 // populate seed entity with minimal data
 businessSeed.setKey("1B5F411F-AAAA-4BF1-B5AB-5FC7555FABA3");
 businessSeed.setName("business entity test 2");
 businessSeed.setOperatorName("myOperatorName");
 businessSeed.setOwnerName("UNAUTHENTICATED");

 // connect to UDDI database and write the minimal entity data
 businessSeed.createStub();

 // create minimal entity object (service)
 ServiceStub serviceSeed = new ServiceStub(keyDetector);

 // populate seed entity with minimal data
 serviceSeed.setKey("3F5F411F-AAAA-4BF1-B5AB-5FC7555FABA3");
 serviceSeed.setName("test service stub");
 serviceSeed.setParentKey("1B5F411F-AAAA-4BF1-B5AB-5FC7555FABA3");
 serviceSeed.setSeq(1);

 // connect to UDDI database and write the minimal entity data
 serviceSeed.createStub();

 System.out.println("Finished writing minimal entities to UDDI registry");

 // this will be thrown if the entity already exists
 } catch (PromoterDBException e) {
 System.out.println(e);
 } catch (PromoterException e) {
 System.out.println(e);
 } catch (DBException e) {
 System.out.println(e);
 }
 }
}

```

```
}
}
}
```

---

## SOAP application programming interface for the UDDI Registry

Access to the SOAP API will by default be available at:

<http://localhost:9080/uddisoap/inquiryapi>

or

<https://localhost:9443/uddisoap/publishapi>

Where 'localhost' is the address by which your WebSphere server is known. If security is enabled on your WebSphere server, the publishapi will also be protected by basic-authentication. By default, when security is enabled, the publishapi is restricted to HTTPS, this is to ensure the confidentiality and security of your data whilst in transit to UDDI. If you do not wish to use SSL, when security is enabled, modify the jar file using AAT Assembly Toolkit to remove the CONFIDENTIAL restriction placed upon the publish URLs. For more information about this topic, see the section on **Configuring SOAP properties with the Application Assembly Tool** **Assembly Toolkit**. If you normally access your WebSphere server via a Web server, ensure the plugin configuration for the WebSphere plugin on the Web server has been updated since installing UDDI. This allows access to the UDDI SOAP API through the URLs :

<http://localhost/uddisoap/inquiryapi>

or

<https://localhost/uddisoap/publishapi>

Where 'localhost' is the address by which your Web server is accessed. Note that if you plan on accessing UDDI via a Web server in this manner, that the samples will require modification to inform them of the SSL certificates used by your Web server, so that the samples can make SSL connections to the Web server. It is beyond the scope of this document to cover the many variants available on Web server/WebSphere/java SSL configurations

- Using the SOAP API
- Handling Errors as a User of the SOAP API

## Programming the SOAP API

To use the SOAP API construct a properly formed UDDI message within the body of a SOAP request, and send it using HTTP POST to the URL of the API that the request relates to. The response is returned within the body of the HTTP reply. Although the samples are written in Java, you can use other programming languages to create your SOAP client, providing you still send requests compliant to the SOAP specification. Valid UDDI requests should conform to the UDDI schema, and be as detailed within the UDDI standard documentation:

<http://www.uddi.org/>

For more information on using the SOAP API, refer to "The UDDI Registry application programming interface".



## SOAP API error handling tips in the UDDI Registry

When using the SOAP API there are three main categories that can cause an error to be returned:

- An incorrect request being sent to the SOAP API., for example:
  - incorrectly formed XML
  - badly formed UDDI requests
  - non-schema compliant requests
- Incorrect business logic within a SOAP API request, for example attempting to delete a business that does not exist.
- Problems occurring while processing a valid request., for example server connection to database failure.

In each of these cases, an error is returned to the client that made the request, which attempts to explain further what the problem was.

---

## UDDI Registry Application Programming Interface

The IBM WebSphere UDDI Registry fully supports the application programming interface (API) specification, which can be viewed by visiting [http://www.uddi.org/pubs/ProgrammersAPI\\_v2.pdf](http://www.uddi.org/pubs/ProgrammersAPI_v2.pdf). Any changes from this specification are documented within the IBM WebSphere UDDI Registry information.

- The Inquiry API
- The Publish API

### Inquiry API for the UDDI Registry

The Inquiry API provides four forms of query that follow broadly used conventions that match the needs of software traditionally used within registries.

- The browse pattern
- The drill-down pattern
- The invocation pattern
- Inquiry API functions

### Browse pattern for the UDDI Registry

Software that allows people to explore and examine data - especially hierarchical data - requires browse capabilities. The browse pattern characteristically involves starting with some broad information, performing a search, finding general result sets and then selecting more specific information for drill-down.

The UDDI API specifications accommodate the browse pattern by way of the *find\_xx* API calls. These calls form the search capabilities provided by the API and are matched with summary return messages that return overview information about the registered information that is associated with the inquiry message type and the search criteria specified in the inquiry.

A typical browse sequence might involve finding whether a particular business you know about has any information registered. This sequence would start with a call to *find\_business*, perhaps passing the first few characters of a business name that you already know. This returns a *businessList* result. This result is overview information (keys, names and descriptions) derived from the registered *businessEntity* information, matching on the name fragment that you provided. If you spot the business you are looking for within this list, you can drill down into the corresponding *businessService* information, looking for particular technical models (for example purchasing, shipping, and so on) using the *find\_service* API call. Similarly, if you know the technical *fingerprint* (tModel signature) of a particular software interface and want to see if the business you have chosen provides a Web service that supports that interface, you can use the *find\_binding* inquiry message.

## Drilldown pattern for the UDDI Registry

When you have a key for one of the four main data types managed by a UDDI registry, you can use that key to access the full registered details for a specific data instance. The UDDI data types are `businessEntity`, `businessService`, `bindingTemplate` and `tModel`. You can access the full registered information for any of these structures by passing a relevant key type to one of the `get_xx` API calls.

Continuing the example from the Browse pattern for the UDDI Registry, one of the data items returned by all of the `find_x` return sets is key information. In the case of the business we were interested in, the `businessKey` value returned within the contents of a `businessList` structure can be passed as an argument to `get_businessDetail`. The successful return to this message is a `businessDetail` message containing the full registered information for the entity whose key value was passed. This will be a full `businessEntity` structure.

## Invocation pattern for the UDDI Registry

To prepare an application to take advantage of a remote Web service that is registered within the UDDI registry by other businesses or entities, you must prepare that application to use the information found in the registry for the specific service being invoked.

The `bindingTemplate` data obtained from the UDDI registry represents the specific details about an instance of a given interface type, including the location at which a program starts interacting with the service. The calling application or program should cache this information and use it to contact the service at the registered address whenever the calling application needs to communicate with the service instance. In previously popular remote procedure technologies tools have automated the tasks associated with caching (or hard coding) location information. Problems arise however when a remote service is moved without any knowledge on the part of the callers. Moves occur for a variety of reasons, including server upgrades, disaster recovery, and service acquisition and business name changes.

When a call fails using cached information previously obtained from a UDDI Registry, the proper behavior is to query the UDDI Registry for fresh `bindingTemplate` information. If the data returned is different from the cached information, the service invocation should automatically retry the invocation using the fresh information. If the result of this retry is successful, the new information should replace the cached information.

By using this pattern with Web services, a business using a UDDI Registry can automate the recovery of a large number of partners without undue communication and coordination costs. For example, if a business has activated a disaster recovery site, most of the calls from partners fail when they try to invoke services at the failed site. By updating the UDDI information with the new address for the service, partners who use the invocation pattern automatically locate the new service information and recover without further administrative action.

## Inquiry API functions in the UDDI Registry

These messages represent inquiries that can be made of the UDDI Registry. These messages all behave synchronously.

The queries available are:

### **find\_binding**

Locates specific bindings within a registered `businessService`. Returns a `bindingDetail` message that contains zero or more `bindingTemplate` structures matching the criteria specified in the argument list.

### **find\_business**

Locates information about one or more businesses. Returns a `businessList` message that matches the conditions specified in the arguments.

**find\_relatedBusinesses**

Locates information about businessEntity registrations that are related to a specific business entity whose key is passed in the inquiry. The Related Businesses feature is used to manage registration of business units and subsequently relate them based on organizational hierarchies or business partner relationships. Returns a relatedBusinessList message containing results that match the conditions specified in the arguments.

**find\_service**

Locates specific services within a registered businessEntity. Returns a serviceList message that matches the conditions specified in the arguments.

**find\_tModel**

Locates a list of tModels that match a set of specified criteria. The response will be a list of abbreviated information about registered tModel data that matches the criteria specified. The result will be returned in a tModelList message.

**get\_bindingDetail**

Requests the run-time bindingTemplate information for the purpose of invoking a registered business API. Returns a bindingDetail message.

**get\_businessDetail**

Returns complete businessEntity information for one or more specified businessEntity registrations matching on the businessKey values specified. Returns a businessDetail message.

**get\_businessDetailExt**

Returns extended businessEntity information for one or more specified businessEntity registrations. This message returns exactly the same information as the get\_businessDetail message, but may contain additional attributes if the source is an external registry with the API specification.

**get\_serviceDetail**

Requests full information about a known businessService structure. Returns a serviceDetail message.

**get\_tModelDetail**

Gets full details for a given set of registered tModel data. Returns a tModelDetail message.

## Publish API for the UDDI Registry

The messages in this section represent commands that are used to publish, delete and update information contained in a UDDI registry. The messages defined in this section all behave synchronously.

The Publishing API calls defined that UDDI operators support are:

**add\_publisherAssertions**

Causes one or more publisherAssertions to be added to an individual publisher's assertion collection.

**delete\_binding**

Causes one or more instances of bindingTemplate data to be deleted from the UDDI registry.

**delete\_business**

Removes one or more business registrations and all direct contents from a UDDI registry.

**delete\_publisherAssertions**

Causes one or more publisherAssertion elements to be removed from a publisher's assertion collection.

**delete\_service**

Removes one or more businessService elements from the UDDI registry and from its containing businessEntity parent.

**delete\_tModel**

Logically deletes one or more tModel structures. Logical deletion hides the deleted tModels from find\_tModel result sets but does not physically delete them, so they are returned on a get\_registeredInfo request.

**discard\_authToken**

Informs an operator site that the authentication token is to be discarded, effectively ending the

session. Subsequent calls that use the same authToken will be rejected. This message is optional for operator sites that do not manage session state or that do not support the get\_authToken message.

#### **get\_assertionStatusReport**

Provides administrative support for determining the status of current and outstanding publisher assertions that involve any of the business registrations managed by the individual publisher account. Using this message, a publisher can see the status of assertions that they have made, as well as see assertions that others have made that involve businessEntity structures controlled by the calling publisher account.

#### **get\_authToken**

Obtains an authentication token. Authentication tokens are opaque values that are required for all other publisher API calls. This message is not required for operator sites that have an external mechanism defined for users to get an authentication token. This API is provided for implementations that do not have some other method of obtaining an authentication token or certificate, or that choose to use userid and password based authentication.

#### **get\_publisherAssertions**

Obtains the full set of publisher assertions that are associated with an individual publisher account. Publisher assertions are used to control publicly visible business relationships.

#### **get\_registeredInfo**

Gets an abbreviated list of all businessEntity and tModel data that are controlled by the individual associated with the credentials passed.

#### **save\_binding**

Saves or updates a complete bindingTemplate element. This message can be used to add or update one or more bindingTemplate elements as well as the container/contained relationship that each bindingTemplate has with one or more existing businessService elements.

#### **save\_business**

Saves or updates information about a complete businessEntity element. This API has the broadest scope of all the save\_xx API calls in the publisher API, and can be used to make sweeping changes to the published information for one or more businessEntity elements controlled by an individual.

#### **save\_service**

Adds or updates one or more businessService elements exposed by a specified businessEntity.

#### **save\_tModel**

Adds or updates one or more registered tModel elements.

#### **set\_publisherAssertions**

Manages all of the tracked relationship assertions associated with an individual publisher account.

For full details of the syntax of the above queries, refer to the API specification at [http://www.uddi.org/pubs/ProgrammersAPI\\_v2.pdf](http://www.uddi.org/pubs/ProgrammersAPI_v2.pdf).

---

## **UDDI EJB Interface for the UDDI Registry**

This section describes how to use the EJB application programming interface (API) of the IBM WebSphere UDDI Registry component to publish, find and delete UDDI entries.

The necessary client classes are contained in the *uddiejbclient.jar* file in the *ejb* subdirectory of the UDDIReg directory under the WebSphere application server directory tree.

The Javadoc for the EJB API is contained in the javadoc directory tree under the *ejb* subdirectory of the UDDIReg directory under the WebSphere appserver directory tree.

The EJB API is contained in two stateless session beans, one for the Inquiry API (*com.ibm.uddi.ejb.InquiryBean*) and one for the Publish API (*com.ibm.uddi.ejb.PublishBean*), whose public methods form an EJB interface for the UDDI Registry. All the public methods on the *InquiryBean* correspond to UDDI Inquiry API functions, and all the public methods on the *PublishBean* correspond to

UDDI Publish API functions. (Not all UDDI API functions are implemented, for example get\_authToken, discard\_authToken, get\_businessDetailExt, and so on) For Version 1 of the UDDI registry, the EJB component supports only UDDI v2.0.

The two EJBs use container-managed transactions. The transaction attribute for the methods of the InquiryBean is NotSupported, and for the methods of the PublishBean it is Required. You must not change the transaction attributes as this could result in undesirable behavior.

Within each interface there are groups of overloaded methods that correspond to the operations in the UDDI 2.0 specification. There is a separate method for each major variation in function. For example, the single UDDI 2.0 operation find\_business is represented by 10 variations of findBusiness methods, with different variations for finding by name, finding by categoryBag and so on.

The arguments for the EJB interface methods are java objects in the package com.ibm.uddi.datatypes. Roughly speaking, there is a one-one correspondence between classes in this package and elements of the UDDI V2.0 XML schema. Exceptions to this are, for example, where UDDI XML elements can be represented by a single String. (See Package com.ibm.uddi.datatypes below for more information.)

### Enabling an EJB Client

This section is written on the assumption that WebSphere Application Server V5.0, a supported database and the IBM WebSphere UDDI Registry have already been installed.

#### Classpaths

Add the following jar files and folders to your CLASSPATH:

##### For Windows

<WebSphere-install-dir>\lib\j2ee.jar
<WebSphere-install-dir>\lib\naming.jar
<WebSphere-install-dir>\lib\namingclient.jar
<WebSphere-install-dir>\lib\ecutils.jar
<WebSphere-install-dir>\lib\sas.jar
<WebSphere-install-dir>\properties

##### For Unix Platforms, and also including z/OS

<WebSphere-install-dir>/lib/j2ee.jar
<WebSphere-install-dir>/lib/naming.jar
<WebSphere-install-dir>/lib/namingclient.jar
<WebSphere-install-dir>/lib/ecutils.jar
<WebSphere-install-dir>/lib/sas.jar
<WebSphere-install-dir>/properties

In addition to these jars, there is also the jar file that contains all of the UDDI specific API for the EJB interface, which can be found at:

##### For Windows

<DeploymentManager-install-dir>\UDDIReg\ejb\uddiejbclient.jar
---------------------------------------------------------------

where <DeploymentManager-install-dir> is the install location for WebSphere Application Server for Network Deployment, which by default is C:\Progra~1\WebSphere\DeploymentManager.

### For Unix Platforms, and also including z/OS

```
<DeploymentManager-install-dir>/UDDIReg/ejb/uddiejbclient.jar
```

where <DeploymentManager-install-dir> is the install location for WebSphere Application Server for Network Deployment, which by default is /opt/WebSphere/DeploymentManager for Linux/Solaris systems or /usr/WebSphere/DeploymentManager for AIX systems.

### The Path

Ensure that your PATH statement starts with <WebSphere-install-dir>\java\bin

### Creating an EJB Client

If you want to read about creating EJB Clients in more detail, then please read the "Sun Microsystems Enterprise JavaBeans™ Specification Version 2.0"

### Finding the EJB Reference

An EJB Client can be a stand-alone Java application, an applet, servlet or a JSP. This document only covers writing a stand-alone Java application. In order to invoke an enterprise java bean (EJB) that has been deployed into WebSphere on the server side, the Client must do two things: find the EJB on the server, and then create a Client side reference to that EJB. When this Client side reference has been created, the Client can invoke methods upon the EJB as if it was a local object. Clients cannot reference, or invoke, and EJB directly. Any calls made to the EJB must be made through the interfaces that the EJB provides. The interface that is used to create a local reference to the EJB is called the *home interface*. When an EJB is deployed in WebSphere, this home interface is made available to Clients by means of a searchable namespace. This means that a Client can look up an address on the namespace. If there is a home interface at that address, and it is the home interface to the EJB that they were looking for, then the Client can create a local instance of that home interface, and then, from that, a local reference to the EJB can be created.

### What code is needed in the Client?

The following code fragment illustrates how to Find and Create a local instance of the Inquiry EJB only. The same must be done to Find and Create a local copy of the Publish EJB.

```
private com.ibm.uddi.ejb.Inquiry inquiry = null;
// This private variable, "inquiry" is going to be the local reference to the EJB in WebSphere
// declaring it outside the scope of a method means that this same reference can be
// used throughout the client, without having to query the namespace again.
public void homeLookup()
{
 // These variables simply determine the address of the JNDI namespace, and the address
 // of the home interface within that namespace.

 // String naming_factory = "com.ibm.ejs.ns.jndi.CNInitialContextFactory";
 //WAS 4.0.2 Naming Factory
 String naming_factory = "com.ibm.websphere.naming.WsnInitialContextFactory";
 //WAS 5.0 Naming Factory

 String namespace_address = "iiop://localhost:2809/"; //The address of the namespace
 String home_address = "com/ibm/uddi/ejb/InquiryHome";
 //The address of the home interface within the JNDI namespace

 java.util.Hashtable environment = new java.util.Hashtable();
```

```

environment.put(javax.naming.Context.INITIAL_CONTEXT_FACTORY, naming_factory);
environment.put(javax.naming.Context.PROVIDER_URL, namespace_address];

try
{
 javax.naming.InitialContext ic = new javax.naming.InitialContext(environment);
 // Create a context using the details above to connect to the namespace

 Object o = ic.lookup(home_address);
 // Do a lookup to see if there is an ejb_home at the address specified above

 // Now create a valid home instance for the EJB type we want to create
 com.ibm.uddi.ejb.InquiryHome home =
 (com.ibm.uddi.ejb.InquiryHome)(javax.rmi.PortableRemoteObject.narrow(o,
 com.ibm.uddi.ejb.InquiryHome.class));

 inquiry = home.create();
 // Now create a local reference of the EJB, by using the home.create() method.
 // Any business method that is intended for the EJB in WebSphere
 // must me invoked against this inquiry object.
}
catch (javax.naming.NamingException ne) {ne.printStackTrace();}
// This is thrown if there was a problem connecting to the namespace, or finding the
// home_address in the namespace
catch (java.rmi.RemoteException re) {re.printStackTrace();}
// This usually indicates some sort of system failure, either WebSphere is not running,
// or there is a communications problem
catch (javax.ejb.CreateException ce) {ce.printStackTrace();}
// This is thrown if the EJB reference cannot be created from the home instance.
}

```

## Writing Client code to use the EJB API

When the reference to the EJB has been created (the Inquiry Object, in the code shown in the previous paragraph), then the reference can be treated like any other Java object. This is an example method using the UDDI EJB API - the only important point to remember is that, although the Inquiry Object has been created as a local reference, it is still referring to a remote EJB Object in a different server, possibly even in a different country. This means that at the very least a `javax.rmi.RemoteException` must be caught on each method call that is made to the EJB.

```

public void findBusiness()
{
 System.out.println("Find Business:");
 NameList names = new NameList();
 names.add(new Name("IBM Corporation"));
 //Create the list of names to find in the UDDI Registry, here just one is used,
 // "IBM Corporation"
 try
 {
 BusinessList list = inquiry.findBusiness(names);
 //This is the call to the inquiry EJB that searches through the UDDI Registry

 //Now display the amount of business found, and for each one, get the BusinessKey,
 // the BusinessName and the amount of Services that Business has
 System.out.println(
 "There are "+list.getBusinessInfos().size()+" matching Businesses in this registry");
 for (int i=0;i<list.getBusinessInfos().size();i++)
 {
 BusinessInfo business = list.getBusinessInfos().get(i);
 System.out.println("\nBusinessKey = "+business.getBusinessKey());
 System.out.println("BusinessName = "+business.getNames().get(0).getNameString());
 System.out.println(
 "This Business Has "+business.getServiceInfos().size()+" Services\n");
 }
 }
}

```



```

// This is a UDDI specific exception, and will be thrown if for example an invalid name
// was used as the search criteria
catch (com.ibm.uddi.datatypes.DispositionReportException e) {
 this.handleDispositionReportException(e);}

 catch (java.rmi.RemoteException re) {re.printStackTrace();}
// This is the RemoteException that is thrown if there has been a system failure or a
// connection problem.
}

```

### What new code is needed on the client?

Just as each EJB has an interface listed on the JNDI namespace, the `javax.transaction.UserTransaction` class also has an interface listed. This means that the same method used to get a local instance of an EJB can be applied to get a local instance of the `UserTransaction` class. Again, this code can be used to find the `UserTransaction` reference on the namespace, in addition to the code required to find the `Inquiry` EJB and the `Publish` EJB, or, alternatively, there is a slightly more elegant method used in the `TransactionEJBClientSample.java`.

```

public void txLookup()
{
 private javax.transaction.UserTransaction tx = null;
 // This is the private variable that will be used to hold the UserTransaction Object
 // declaring it outside the scope of a method means that this same reference can be
 // used throughout the client, without having to query the namespace again.

 // These variables simply determine the address of the JNDI namespace, and the address
 // of the home interface within that namespace.

 // String naming_factory = "com.ibm.ejs.ns.jndi.CNInitialContextFactory";
 //WAS 4.0.2 Naming Factory
 String naming_factory = "com.ibm.websphere.naming.WsnInitialContextFactory";
 //WAS 5.0 Naming Factory

 String namespace_address = "iiop://localhost:2809/"; //The address of the namespace
 String transaction_address = "jta/usertransaction";
 //The address of the UserTransaction interface within the JNDI namespace

 java.util.Hashtable environment = new java.util.Hashtable();
 environment.put(javax.naming.Context.INITIAL_CONTEXT_FACTORY, naming_factory);
 environment.put(javax.naming.Context.PROVIDER_URL, namespace_address);

 try
 {
 javax.naming.InitialContext ic = new javax.naming.InitialContext(environment);
 // Create a context using the details above to connect to the namespace
 Object remote_object = ic.lookup(transaction_address);
 // Do a lookup to see if there is a UserTransaction Object at the address specified above
 tx = (javax.transaction.UserTransaction)remote_object;
 //Convert the remote object found into a UserTransaction Object, and assign to the private variable
 }
 catch (javax.naming.NamingException ne) {ne.printStackTrace();}
 // This is thrown if there was a problem connecting to the namespace, or finding the
 // transaction_address in the namespace
}

```

### Writing Client code to use the EJB API with a Client transaction

To perform an `Inquiry`, a `Publish` or a `Delete` upon the IBM WebSphere UDDI Registry with client-side transactional support requires very little additional code compared to doing the same operations without client side transactional support. Using the same code that is listed previously (in "Writing Client Code to use the EJB API"), this example illustrates how easy client side transactions are to implement.

The additional lines of code needed are in bold type. This code also assumes that there is a variable called *tx* that has been declared at the class scope.

```
public void findBusiness()
{
//Just as there are UDDI and RMI specific exceptions thrown, 5 more exceptions need to be caught.
try
{
tx.begin(); //This begins the transaction context
System.out.println("Find Business:");
NameList names = new NameList();
names.add(new Name("IBM Corporation"));
//Create the list of names to find in the UDDI Registry, here just one is used, "IBM Corporation"
try
{
BusinessList list = inquiry.findBusiness(names);
//This is the call to the inquiry EJB that searches through the UDDI Registry

//Now display the amount of business found, and for each one, get the BusinessKey,
// the BusinessName and the amount of Services that Business has
System.out.println("There are "+list.getBusinessInfos().size()+" matching Businesses in this registry");
for (int i=0;i<list.getBusinessInfos().size();i++)
{
BusinessInfo business = list.getBusinessInfos().get(i);
System.out.println("\nBusinessKey = "+business.getBusinessKey());
System.out.println("BusinessName = "+business.getNames().get(0).getNameString());
System.out.println("This Business Has "+business.getServiceInfos().size()+" Services\n");
}
}
// This is a UDDI specific exception, and will be thrown if for example an invalid name was
// used as the search criteria
catch (com.ibm.uddi.datatypes.DispositionReportException e) {
this.handleDispositionReportException(e);}
catch (java.rmi.RemoteException re) {re.printStackTrace();}
// This is the RemoteException that is thrown if there has been a system failure or a
// connection problem.

tx.commit(); //This ends the transaction context
}
catch (javax.transaction.NotSupportedException nse) {nse.printStackTrace();}
catch (javax.transaction.RollbackException rbe) {rbe.printStackTrace();}
catch (javax.transaction.SystemException se) {se.printStackTrace();}
catch (javax.transaction.HeuristicMixedException hme) {hme.printStackTrace();}
catch (javax.transaction.HeuristicRollbackException hrbe) {hrbe.printStackTrace();}
}

```

- The datatypes package
- Methods in the EJB Interface

## Datatypes package in the UDDI Registry

The following table lists the classes in the com.ibm.uddi.datatypes package, the elements in the UDDI v2.0 XML schema, and the correspondence between the two.

com.ibm.uddi.datatypes Class	Corresponding UDDIv2.0 XML Schema Element	Notes on DatatypeClass
AccessPoint	accessPoint	
Address	address	
	String addressLine	
AdressLineList		Encapsulates a vector of addressLine Strings
AddressList		Encapsulates a vector of Address objects
AssertionStatusItem	assertionStatusItem	

AssertionStatusItemList		Encapsulates a vector of AssertionStatusItem objects
AssertionStatusReport	assertionStatusReport (response message)	
	String authInfo	
AuthToken		Object containing authInfo String and operator name
	String bindingKey	
BindingDetail	bindingDetail (response message)	
BindingTemplate	bindingTemplate	
BindingTemplateList	bindingTemplates	Encapsulates a vector of Bindingtemplate objects
BusinessDetail	businessDetail (response message)	
BusinessDetailExt	businessDetailExt (Response message)	**
BusinessEntity	businessEntity	
BusinessEntityExt	businessEntityExt	**
BusinessEntityExtList		Encapsulates a vector of BusinessEntityExt objects **
BusinessEntityList		Encapsulates a vector of BusinessEntity objects
BusinessInfo	businessInfo	
BusinessInfoList	businessInfo	Encapsulates a vector of businessInfo objects
	String businessKey	
BusinessList	businessList (response message)	
BusinessService	businessService	
BusinessServiceList	businessServices	Encapsulates a Vector of BusinessService objects
CategoryBag	categoryBag	
	String completionStatus	
Contact	contact	
ContactList	contacts	Encapsulates a vector of Contact objects
Description	description	
DescriptionList		Encapsulates a vector of Description objects
DiscoveryUrl	discoveryURL	
DiscoveryUrlList	discoveryURLs	Encapsulates a vector of DiscoveryURL objects
DispositionReport	dispositionReport	
DispositionreportException		Exception thrown by EJB interface functions when an error occurs
Email	email	
EmailList		Encapsulates a vector of Email objects
EndPoint		Used as baseclass for AccessPoint and HostingRedirector providing mutual exclusivity
ErrInfo	errInfo	
	findQualifier	
FindQualifier	findQualifiers	
	String fromKey	

HostingRedirector	hostingRedirector	
IdentifierBag	identifierbag	
InquiryOptions		Encapsulates a FindQualifiers object and a maxrows field. Used in find_* API calls to specify search options
InstanceDetails	instanceDetails	
	String instanceParms	
	String keyValue	
KeyedReference	keyedReference	
keysOwned	keysOwned	
LanguageString		Abstract class, extended by some of the datatypes, which represents a string that can optionally be tagged with xml:lang.
Name	name	
NameList		Encapsulates a vector of Name objects
OverviewDoc	overviewDoc	
	String overviewURL	
	String personName	
Phone	phone	
PhoneList		Encapsulates a vector of Phone objects
PublisherAssertion	publisherAssertion	
PublisherAssertionList		Encapsulates a vector of Publisher Assertion objects
PublisherAssertions	publisherAssertions (response message)	
RegisteredInfo	registeredInfo (response message)	
	relatedBusinessInfo	Not used
	relatedBusinessInfos	Not used
RelatesBusinessesList	relatedBusinessesList	
RelatedBusinessInfo	relatedBusinessInfo	
RelatedBusinessInfos	relatedBusinessInfos	
Result	result	
ResultList		Encapsulates a Vector of Result objects
ServiceDetail	serviceDetail (response message)	
ServiceInfo	serviceInfo	
ServiceInfoList	serviceInfos	Encapsulates a vector of serviceInfo objects
	String serviceKey	
ServiceList	serviceList (response message)	
	sharedRelationships	Not used
SharedRelationships	sharedRelationships	
Tmodel	tModel	
TModelBag	tModelBag	

TModelDetail	tModelDetail (response message)	
TModelInfo	tModelInfo	
TModelInfoList	tModelInfos	Encapsulates a vector of TModelInfo objects
TModelInstanceInfo	tModelInstanceInfo	
TModelInstanceInfoList	tModelInstanceDetails	Encapsulates a vector of TModelInstanceInfo objects
	String tModelKey	
TModelList	tModelList (response message)	
TModels		Encapsulates a vector of TModel objects
	String toKey	
	String uploadRegister	
UploadRegisterList		Encapsulates a vector of uploadRegister strings

\*\* Used in UDDI API functions relating to BusinessDetailExtension. These UDDI API functions are not implemented in Version 1 of the IBM WebSphere UDDI Registry.

In general, a datatype called *DatatypeList* contains a vector of *Datatype* objects. Often these correspond to XML schema elements with plural names. (For example the datatype *Contact* corresponds to XML element *contact*, and *ContactList* corresponds to *contacts*.) Where there is no "plural" XML schema element for a particular *Datatype*, often there is still a *DatatypeList* where it is useful to have one, for example *AddressList*.

The exceptions to this naming convention occur when there is an existing XML schema element ending in "List". The exceptions are: *TModelList*, *ServiceList*, *BusinessList*. In these cases, the corresponding datatypes are given the same names as the XML schema elements, and the datatypes that would have had these names are called: *TModels*, *BusinessServiceList*, *BusinessEntityList*.

## EJB interface methods in the UDDI Registry

### Inquiry

```
findBinding
findBusiness
findRelatedBusinesses
findService
findTModel
getBindingDetail
getBusinessDetail
getServiceDetail
getTModelDetail
```

### Publish

```
addPublisherAssertions
deleteBinding
deleteBusiness
deletePublisherAssertions
deleteService
deleteTModel
getAssertionStatusReport
getRegisteredInfo
getPublisherAssertions
saveBinding
```

```
saveBusiness
saveService
saveTModel
setPublisherAssertions
```

Each method is overloaded and can take various combinations of arguments. The Javadoc information contains detailed information about each method.

Note that the `get_authToken` and `discard_authToken` methods are not implemented, as WebSphere security is used instead.

---

## UDDI troubleshooting tips

When the IBM WebSphere UDDI Registry is running, it might issue messages to report events or errors. You can use these messages, described in Messages as your first aid to problem determination. If you need more details about the causes of a problem, you can turn on tracing for UDDI, as described in:

- Turning on UDDI trace

- **Common causes of errors**

Below are a few of the common causes of errors that might be found and their suggested solutions.

- If you set up the UDDI Registry application with a JDBC driver and datasource that reference Cloudscape, but set the `persist` property in `uddi.properties` to specify DB2, **or vice versa**, then some unexpected behavior results, such as a fatal error on deleting an entity. If this happens, you should check that the above details are not in conflict. This only applies to a UDDI Registry installation on a single appserver.
- If you get a message "The application failed to initialize" when trying to access the UDDI User console and you are using DB2 as the persistence store for the UDDI Registry, a likely cause of the problem is that you specified the wrong userid and/or password when you ran the script to install the UDDI Registry application. If this occurs rerun the script ensuring you use the correct userid and password.

Alternatively, on Unix platforms, you may not have run the `db2profile` before deploying the UDDIReg application or before starting the application server.

- You might find that, after uninstalling and reinstalling the UDDI Registry, you get errors from the UDDI User Console of the form:

```
"Error 500: JSPG0059E: Unable to compile class for JSP".
```

If this occurs, then you should clear out the temp directory of the WebSphere AppServer.

- When running one of the UDDI setup scripts `setupuddi.jacl` or `removeuddi.jacl`, if you get an error such as:

```
WASX7017E: Exception received while running file "setupuddi.jacl";
exception information: com.ibm.bsf.BSFException:
error while eval'ing Jacl expression:
java.util.MissingResourceException:
```

```
Can't find resource for bundle java.util.PropertyResourceBundle, key ErrMsgIncorrectNumArgs
```

ensure that the file `setupuddimessages.jar` is located in the `lib` subdirectory of the WebSphere deployment manager or application server under which you are running the script.

- When running the DB2 Setup Wizard, if you get an error stating "Invalid userid and password", it could be caused by any of the following situations:
  - You have supplied an invalid userid or password - re-enter with a valid userid and password.
  - The supplied userid does not have the necessary privileges - retry with a userid that has appropriate privileges.
  - DB2 is stopped when you run the wizard - start DB2 and retry the wizard.
  - The UDDI20 database already exists and has been removed previously and, as such, is not catalogued. The DB2 wizard does not recognize this situation and gives the error. You now have two options.
    1. To use the existing database, catalogue it, and there is no need to rerun the Wizard.

2. To create a new database, recatalog the database and re-run the DB2 wizard and choose the option to overwrite the database. (Any existing data **WILL** be lost.)
  - Catalog the database by:

- **On Windows:**

```
>db2cmd
>db2cat -d uddi20
```

- **On Unix platforms**

```
>su db2inst1 (or name of your db2 instance)
>db2 CATALOG DATABASE UDDI20
```

- ensure that, if you are using a non-English installation of DB2, you have applied PTF501.

**Note:** If you have a copy of the file SetupDB2UDDI.jar in your appserver directory, the application of the base and Network Deployment PTFs will not update SetupDB2UDDI.jar in your appserver directory. You must apply the PTF for Network Deployment to your DeploymentManager file structure to update the SetupDB2UDDI.jar located there (in the /IDDIReg/scripts subdirectory), and then manually copy this jar file to any application server you may wish to run it on.

- There is a limitation concerning URL rewriting causing JavaScript syntax errors on several Web pages in the UDDI User Console. Because of this, cookies must be enabled in client browsers, the application server must have cookies enabled as the session tracking mechanism, and URL rewriting must be disabled.
- If you have an existing DB2 version of the UDDI Registry database, and you use the UDDI DB2 setup wizard to replace this database with a new one, and if the database is in use at the time that you run the UDDI DB2 setup wizard, then the existing database is not overwritten.
- When running the UDDI DB2 setup wizard, as part of the installation step "Setting up the UDDI Registry to use DB2 within a deployment manager cell" or "Setting up the UDDI Registry to use DB2 in a single application server", in addition to running the *was\_install\bin\setupcmdline.bat* directory, you should also enter either **set PATH=%WAS\_PATH%** (for Windows platforms) or **export PATH=/opt/WebSphere/AppServer/java/bin:\$PATH** (for Unix platforms) to ensure that you have access to Java.
- UDDI user console "Page cannot be displayed" errors with Internet Explorer.
 

If you use Internet Explorer with the option "Show friendly HTTP error messages" enabled and you have WebSphere Application Server security enabled (user ID and password authentication enabled), you might experience intermittent errors on the browser, such as "Page cannot be displayed", when navigating the UDDI user console. This might be particularly noticeable when accessing the publish actions.

To avoid such errors, disable the "Show friendly HTTP error messages" option on Internet Explorer. This option is found under **Tools > Internet options > Advanced Tab > Browsing Section**
- When using SOAP or UDDI4J, it is sometimes necessary to call `setServiceKey("")` before saving your changes, except with the EJB interface where this might result in an error.
- There are known problems with inquiries issued against the UDDI Registry if IBM Cloudscape is used as the persistence store for the registry data. Certain complex inquiries might produce unexpected results, or could fail. If your application needs to make inquiries of this nature, consider using DB2 as the persistence store. Note: DB2 must be used for production purposes. The IBM Cloudscape support is only provided for development and test use.
- You might see errors if you specify requests that specify more than 5 category values, more than 5 identifier values, or more than 5 technical model (tModel) values.
- If you stop and restart the UDDI Registry application from the administrative console, and then try to access the Registry through the user console, an "Error 500 - object is not an instance of a declaring class" displays on the user console, and the error message "SRVE0026E" displays in the system log. You cannot access the UDDI Registry until you restart the WebSphere Application Server. To avoid this problem when restarting the UDDI Registry, you should set **Prefer WEB-INF classes** on the panel navigated to by the following steps: **Applications > Manage Applications > UDDIRegistry > Web Modules defined for this Application > gui.war**.
- It is possible that a scripting error displays when you are running the **wsadmin appserversetupuddi.jacl** command.



During installation, if you see the following error at the end of running the **appserversetupuddi.jacl** command, you can safely ignore the error. It is recommended that you start (or stop and restart) the application server and then continue.

Here is an example of the error:

```
UDIN2041I: Starting UDDI application. UDIN8019E: startApplication command for
 appname caught exception Exc. Values are: appname=UDDIRegistry,
 Exc=com.ibm.ws.scripting.Scripting
 Exception: com.ibm.websphere.management.exception.Connector
 Exception: ADMC0009E: Failed to make the SOAP RPC call: invoke
```

- If you find that the UDDI user console is giving unexpected results, in particular if you have installed a FixPack but are not seeing all of the fixes, then you should clear out the compiled UDDI JSP files from the temp directory of the WebSphere Application Server. These files will be in a directory identified by the name of the UDDI Registry located beneath the temp directory.
- If attempting to use a remote DB2 database and you are experiencing problems attaching to the remote system, one of the possible causes might be IP addressing. You should not have this problem if the remote system is using a static IP address. If, however, the remote system is using DHCP, the two systems must be aware of each others subnet mask.

For Windows, the subnet mask can be found by starting a Command Prompt and entering "*ipconfig*" on the remote system. On the host system, the WINS might need to be edited to add the remote subnet mask. To do this on Windows go to the following commands:

1. START => Network and Dial-up Connections => Local Area Network Connection 2 => Internet Protocol (TCP/IP) and click on Properties
2. Click on "*Advanced*".
3. Click on the WINS tab and add the new subnet mask
4. Move the new subnet mask to the top of the list by highlighting it and pressing the "*Up*" arrow until it is the top of the list of WINS addresses

On Unix platforms, you can use *ifconfig* to determine the subnet mask.

- **Known limitations with UDDI Utility Tools and workarounds**

There are known limitations with the UDDI Utility Tools and a workaround for each:

- Referenced businesses in service projections are not added automatically to the EDF in the same manner as referenced tModels.

**Workaround:** Add the referenced business that will 'own' the projected service to the EDF. If the business is not present in the target registry, it should be placed before the service's owning business in the EDF.

- Cycle detection for service projections are not detected in the same manner as for referenced tModels.

**Workaround:** If a circular reference is present between two or more service projections, break the cycle by removing one of the projections temporarily, perform the import and update the changed entity to reestablish the cycle in the target registry.

- tModels that were deleted (in the logical sense) in the source registry are imported and promoted as undeleted in the target registry. This is because, in the UDDI Version 2 specification, the deleted state of tModels is not exposed as API calls.

**Workaround:** After importing the tModel, perform a delete. This is done using the UDDI Utility Tools delete function, or any other UDDI Registry API access method.

- BindingTemplates referenced by hostingRedirectors are not added automatically to the EDF in the same manner as referenced tModels.

**Workaround:** Add the referenced bindingTemplate to the EDF.

- Businesses referenced by an 'owningBusiness' keyedReference are not automatically added to the EDF.

**Workaround:** Import the referenced business into the target registry before importing the tModel that references it.

- The JSSE provider class, when security is enabled, is not configurable. It must be `com.ibm.jsse.IBMJSSEProvider`.
- A few combinations of command line arguments are not validated and prevented, for example, it is possible to specify `-import with -keysFile <path to file>` in the same command, although the `&#8209;keysFile` is ignored.

## Turning on UDDI trace

You enable UDDI-specific trace in the same way as you enable other tracing in the WebSphere Application Server.

The following is a list of trace strings that you can use:

- `com.ibm.uddi.api`
- `com.ibm.uddi.config`
- `com.ibm.uddi.datatypes`
- `com.ibm.uddi.dom`
- `com.ibm.uddi.ejb`
- `com.ibm.uddi.exception`
- `com.ibm.uddi.exceptions`
- `com.ibm.uddi.gui`
- `com.ibm.uddi.gui.inquire`
- `com.ibm.uddi.gui.publish`
- `com.ibm.uddi.persistence`
- `com.ibm.uddi.persistence.jdbc`
- `com.ibm.uddi.persistence.jdbc.cloudscape`
- `com.ibm.uddi.persistence.jdbc.db2`
- `com.ibm.uddi.ras`
- `com.ibm.uddi.security`
- `com.ibm.uddi.soap`
- `com.ibm.uddi.uuid`
- `com.ibm.uddi.validation`
- `com.ibm.uddi.xml`

For example, to trace the UDDI user console:

```
'com.ibm.uddi.gui=all=enabled'
```

This enables all types of trace for the gui. Please refer to "Enabling trace" in the Information Center for more information about using the administrative console to enable or disable trace.

---

## Messages

When the IBM WebSphere UDDI Registry is running, it issues messages to report events or errors. The messages are in the form `UDxxnnnns` where:

- xx** is a two character descriptor identifying which component is involved
- nnnn** is the error code
- s** is either I (Information) or E (Error)

The prefix `UDxxnnnns`: is followed by text that describes the event or error. For some messages, the first word of the text is one of the form `(MSN=SSSS)`. The SSSS provides a message sequence number (or MSN), which identifies the unique circumstance in which the message was issued, and is of use where the same message can be issued in more than one circumstance.

To help you diagnose problems and minimize the need to enable trace in any of the above components, view the messages table. You can view the messages by prefix or component, whichever is easiest for you to find in the table. All messages are documented with user/system action and explanation.

The text for the UDDI messages is contained in a file *uddiresourcebundles.jar* which is placed, by the installation process, into the `\lib` subdirectory (Windows) of the WebSphere application server into which the UDDI Registry was installed. If you will be running a console or log analyzer from another process; for example, to analyze the activity log, you must place a copy of *uddiresourcebundles.jar* into a directory that is within the classpath of that process. Otherwise, the message lookup for the UDDI messages will fail.

UDDI Components Message Prefix Table	
Click on individual links for message documentation for the component	
UDAI	API
UDCF	Configuration
UDDA	Datatypes
UDDM	DOM
UDEJ	EJB Interface
UDEX	Exceptions
UDIN	Installation
UDLC	Local API
UDPR	Persistence
UDRS	Logging
UDSC	Security
UDSP	SOAP Interface
UDUC	User Console
UDUT	UDDI Utility Tool
UDUU	UUID

## UDAI (Web Services UDDI) messages

There are no messages issued by this component.

## UDCF (Web Services UDDI) messages

**UDCF0001E: Exception occurred while getting int value of configuration property "<property>", exception: "<exception>"**

**Explanation:** This message is issued when an attempt to read the value of a configuration property from the `uddi.properties` file and convert it to integer has failed with the indicated exception.

**User Response:** Check that the `uddi.properties` file contains a value for the indicated configuration property, and that the value is valid. Check also that the indicated configuration property is a legal property. Refer to the Information Center for further information about global configuration properties and the `uddi.properties` file.

**UDCF0002E: Exception occurred while getting long value of configuration property "<property>", exception: "<exception>"**

**Explanation:** This message is issued when an attempt to read the value of a configuration property from the `uddi.properties` file and convert it to long has failed with the indicated exception.

**User Response:** Check that the `uddi.properties` file contains a value for the indicated configuration property, and that the value is valid. Check also that the indicated configuration property is a legal property. Refer to the Information Center for further information about global configuration properties and the `uddi.properties` file.

**UDCF0003E: Exception occurred while getting boolean value of configuration property**

"<property>", exception: "<exception>"

**Explanation:** This message is issued when an attempt to read the value of a configuration property from the uddi.properties file and convert it to boolean has failed with the indicated exception

**User Response:** Check that the uddi.properties file contains a value for the indicated configuration property, and that the value is valid. Check also that the indicated configuration property is a legal property. Refer to the Information Center for further information about global configuration properties and the uddi.properties file.

**UDCF0004E: Failed to load UDDI global properties file.**

**Explanation:** This message is issued when the UDDI global configuration properties file, uddi.properties, cannot be loaded. Default values for the global configuration properties will be set, but these defaults may not be adequate for many of the properties, so you should investigate and resolve this problem.

**User Response:** Check that the uddi.properties file exists and is in the correct directory. Refer to the Information Center for further information about global configuration properties and the uddi.properties file.

**UDCF0005E: Exception occurred while loading UDDI global configuration properties, exception:**

"<exception>"

**Explanation:** This message is issued when an attempt to load the UDDI global configuration properties from the uddi.properties has failed with the indicated exception. Default values for the global configuration properties will be set, but these defaults may not be adequate for many of the properties, so you should investigate and resolve this problem.

**User Response:** Check that the uddi.properties file exists and contains valid values for each of the configuration properties. Refer to the Information Center for further information about global configuration properties and the uddi.properties file.

**UDDA (Web Services UDDI) messages**

There are no messages issued by this component.

**UDDM (Web Services UDDI) messages**

There are no messages issued by this component.

**UDEJ (Web Services UDDI) messages**

There are no messages issued by this component.

**UDEX (Web Services UDDI) messages**

There are no messages issued by this component.

**UDIN (Web Services UDDI) messages****UDIN0001I: Assuming hard coded defaults.**

**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.

**User Response:** None..

**UDIN0002I: Cloudscape classpath is clpath. Value is:**

**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.

**User Response:** None.

**UDIN0003I: Looking for childtype childname under parenttype parentname. Values are:**

**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.

**User Response:** None.

**UDIN0004I: Looking for childtype childname under parenttype parentname and parenttype2 parentname2. Values are:**

**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.

**User Response:** None.

**UDIN0005I: Conflict found with existing childtype childname. Values are:**

**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.

**User Response:** None.

**UDIN0006I: Not creating requested childtype. Value is:**

**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.

**User Response:** None.

**UDIN0007I: Seeking parenttype with requested id of parentname. Values are:**

**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.

**User Response:** None.

**UDIN0008I: Seeking parenttype with requested id of parentname under parenttype2 parentname2. Values are:**

**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.

**User Response:** None.

**UDIN0009I: Attempting to create childtype under parenttype of parentID. Values are:**

**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.

**User Response:** None.

**UDIN0010I: Create command that will be issued is:**

**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.

**User Response:** None.

**UDIN0011I: childtype childid was successfully created. Values are:**

**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.

**User Response:** None.

**UDIN0012I: Looking for builtin\_rra.**

**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.

**User Response:** None.

**UDIN0013I: List for J2CResourceAdapter returned N members. Value is:**

**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.

**User Response:** None.

**UDIN0014I: Hunting J2CResourceAdapter associated with Node nodename. Value is:**

**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.

**User Response:** None.

**UDIN0015I: Using rraID as builtin\_rra. Value is:**

**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.

**User Response:** None.

**UDIN0016I: Using provider class of implclass with a classpath of clpath. Values are:**

**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.

**User Response:** None.

**UDIN0017I: Installing to server servename, node nodename using database type of dbtype. Values are:**

**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.

**User Response:** None.

**UDIN0018I: Attempting to create UDDI JDBCProvider.**

**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.

**User Response:** None.

- UDIN0019I: Attempting to create UDDI Datasource.**  
**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.  
**User Response:** None.
- UDIN0020I: Application Manager appmgr found. Value is:**  
**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.  
**User Response:** None.
- UDIN0021I: Attempting to install UDDI Registry application.**  
**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.  
**User Response:** None.
- UDIN0022I: Checking for installed UDDI Registry application of name appname. Value is:**  
**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.  
**User Response:** None.
- UDIN0023W: Application of name appname is not present. Value is:**  
**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.  
**User Response:** None.
- UDIN0024I: ApplicationManager not running, so application will not need to be stopped.**  
**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.  
**User Response:** None.
- UDIN0025I: Stopping application of name appname. Value is:**  
**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.  
**User Response:** None.
- UDIN0026W: stopApplication command for application appname caught exception Exc. Application might not have been running on this server. Values are:**  
**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.  
**User Response:** None.
- UDIN0027I: Application appname stopped successfully. Value is:**  
**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.  
**User Response:** None.
- UDIN0028I: Removing application appname. Value is:**  
**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.  
**User Response:** None.
- UDIN0029I: Application appname removed successfully. Value is:**  
**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.  
**User Response:** None.
- UDIN0030I: Adding resource bundles to repository.**  
**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.  
**User Response:** None.
- UDIN0031I: Adding Cloudscape user functions to repository.**  
**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.  
**User Response:** None.
- UDIN0032I: UDDI configuration properties file already exists. Only the persister and getServletURLPrefix properties will be overwritten.**  
**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.  
**User Response:** None.
- UDIN0033I: Editing UDDI configuration properties file propsfile. Value is:**  
**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.  
**User Response:** None.

**UDIN0034I: Url prefix found. Updating it to discoveryURL. Value is:**  
**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.  
**User Response:** None.

**UDIN0035I: Persister property found. Updating it to dbtype. Value is:**  
**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.  
**User Response:** None.

**UDIN0036I: Adding UDDI configuration properties file to repository for cell cellname under target node and server. Value is:**  
**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.  
**User Response:** None.

**UDIN0037I: ws.ext.dir processing starting.**  
**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.  
**User Response:** None.

**UDIN0038I: serverID is:**  
**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.  
**User Response:** None.

**UDIN0039I: JVM is:**  
**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.  
**User Response:** None.

**UDIN0040I: Out of N properties we located M matches at positions poslist. Values are:**  
**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.  
**User Response:** None.

**UDIN0041I: Building new ws.ext.dirs properties.**  
**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.  
**User Response:** None.

**UDIN0042I: SYSPROP is:**  
**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.  
**User Response:** None.

**UDIN0043I: ws.ext.dir has been set with new sysprop. Value is:**  
**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.  
**User Response:** None.

**UDIN0044I: ws.ext.dir update skipped, required changes already present.**  
**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.  
**User Response:** None.

**UDIN0045I: ws.ext.dir processing step complete.**  
**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.  
**User Response:** None.

**UDIN0046I: Cleaning up temporary version of properties file temppropsfile. Value is:**  
**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.  
**User Response:** None.

**UDIN0047I: Issuing nodeSync.**  
**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.  
**User Response:** None.

**UDIN0048I: UDDI Registry successfully installed. Please restart server servername to activate configuration changes. Value is:**  
**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.  
**User Response:** None.



- UDIN0049I: Application Manager appmgr found. Value is:**  
**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.  
**User Response:** None.
- UDIN0050I: Server is not running, so will not need to be stopped.**  
**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.  
**User Response:** None.
- UDIN0051I: Stopping server servername under node nodename. Values are:**  
**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.  
**User Response:** None.
- UDIN0052I: Server servername stopped successfully. Value is:**  
**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.  
**User Response:** None.
- UDIN0053I: Restarting application server**  
**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.  
**User Response:** None.
- UDIN0054I: Application server servername restarted successfully. Value is:**  
**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.  
**User Response:** None.
- UDIN0055I: Please ignore any errors concerning the serverStartupSyncEnabled attribute.**  
**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.  
**User Response:** None.
- UDIN0101I: Attempting to save new configuration.**  
**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.  
**User Response:** None.
- UDIN0102I: Changes saved successfully.**  
**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.  
**User Response:** None.
- UDIN0103I: Changes were not saved on this call.**  
**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.  
**User Response:** None.
- UDIN0104I: Attempting to save new configuration.**  
**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.  
**User Response:** None.
- UDIN0105I: Changes saved successfully.**  
**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.  
**User Response:** None.
- UDIN0106I: Attempting to save ws.ext.dir changes.**  
**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.  
**User Response:** None.
- UDIN0107I: Changes saved successfully.**  
**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.  
**User Response:** None.
- UDIN0108I: Attempting final save of new configuration.**  
**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.  
**User Response:** None.
- UDIN0109I: Changes saved successfully.**  
**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.  
**User Response:** None.

- UDIN1001I: Application Manager appmgr found. Value is:**  
**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.  
**User Response:** None.
- UDIN1002I: Server is not running, so will not need to be stopped.**  
**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.  
**User Response:** None.
- UDIN1003I: Stopping server servername under node nodename. Values are:**  
**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.  
**User Response:** None.
- UDIN1004I: Server servername stopped successfully. Value is:**  
**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.  
**User Response:** None.
- UDIN1005I: Resource bundles file will be removed from repository if present.**  
**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.  
**User Response:** None.
- UDIN1006I: Removing resource bundles from repository.**  
**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.  
**User Response:** None.
- UDIN1007I: Resource bundles successfully removed from repository.**  
**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.  
**User Response:** None.
- UDIN1008I: Cloudscape user functions file will be removed from repository if present.**  
**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.  
**User Response:** None.
- UDIN1009I: Removing Cloudscape user functions from repository.**  
**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.  
**User Response:** None.
- UDIN1010I: Cloudscape user functions successfully removed from repository.**  
**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.  
**User Response:** None.
- UDIN1011I: Application Manager appmgr found. Value is:**  
**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.  
**User Response:** None.
- UDIN1012I: Checking for installed UDDI Registry application of name appname. Value is:**  
**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.  
**User Response:** None.
- UDIN1013W: Application of name appname is not present. Value is:**  
**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.  
**User Response:** None.
- UDIN1014I: ApplicationManager not running, so application will not need to be stopped.**  
**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.  
**User Response:** None.
- UDIN1015I: Stopping application of name appname. Value is:**  
**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.  
**User Response:** None.
- UDIN1016W: stopApplication command for application appname caught exception Exc. Application might not have been running on this server. Values are:**  
**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.

**User Response:** None.  
**UDIN1017I: Application appname stopped successfully. Value is:**  
**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.

**User Response:** None.  
**UDIN1018I: Removing application appname. Value is:**  
**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.

**User Response:** None.  
**UDIN1019I: Application appname removed successfully. Value is:**  
**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.

**User Response:** None.  
**UDIN1020I: UDDI datasource will be removed from server servername in node nodename if present. Values are:**  
**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.

**User Response:** None.  
**UDIN1021I: Removing UDDI datasource.**  
**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.

**User Response:** None.  
**UDIN1022I: UDDI datasource successfully removed.**  
**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.

**User Response:** None.  
**UDIN1023I: UDDI JDBC driver will be removed from server servername in node nodename if present. Values are:**  
**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.

**User Response:** None.  
**UDIN1024I: Removing UDDI JDBC driver.**  
**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.

**User Response:** None.  
**UDIN1025I: UDDI JDBC driver successfully removed.**  
**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.

**User Response:** None.  
**UDIN1026I: UDDI configuration properties file will be removed from repository if present.**  
**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.

**User Response:** None.  
**UDIN1027I: Removing configuration properties file from cell cellname, node nodename and server servername. Values are:**  
**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.

**User Response:** None.  
**UDIN1028I: Configuration properties file successfully removed from repository.**  
**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.

**User Response:** None.  
**UDIN1029I: Issuing nodeSync.**  
**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.

**User Response:** None.  
**UDIN1030I: UDDI Registry application, JDBC driver and datasource removed successfully.**  
**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.

**User Response:** None.  
**UDIN1031I: Restarting application server.**  
**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.

**User Response:** None.

**UDIN1032I: Application server servername restarted successfully. Value is:**  
**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.

**User Response:** None.

**UDIN1033I: Please ignore any errors concerning the serverStartupSyncEnabled attribute.**  
**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.

**User Response:** None.

**UDIN1034I: ws.ext.dir processing starting.**  
**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.

**User Response:** None.

**UDIN1035I: serverID is:**  
**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.

**User Response:** None.

**UDIN1036I: JVM is:**  
**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.

**User Response:** None.

**UDIN1037I: Out of N properties we located M matches at positions poslist. Values are:**  
**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.

**User Response:** None.

**UDIN1038I: Removing UDDI values from ws.ext.dirs properties.**  
**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.

**User Response:** None.

**UDIN1039I: ws.ext.dir has been set with new sysprop. Value is:**  
**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.

**User Response:** None.

**UDIN1040I: ws.ext.dir update skipped, required changes already present.**  
**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.

**User Response:** None.

**UDIN1041I: ws.ext.dir processing step complete.**  
**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.

**User Response:** None.

**UDIN1101I: Attempting to save new configuration.**  
**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.

**User Response:** None.

**UDIN1102I: Changes saved successfully.**  
**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.

**User Response:** None.

**UDIN1103I: Attempting to save new configuration.**  
**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.

**User Response:** None.

**UDIN1104I: Changes saved successfully.**  
**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.

**User Response:** None.

**UDIN1105I: Attempting to save new configuration.**  
**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.

**User Response:** None.

**UDIN1106I: Changes saved successfully.**  
**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.

**User Response:** None.

**UDIN1107I: Attempting final save of new configuration.**

**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.

**User Response:** None.

**UDIN1108I: Changes saved successfully.**

**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.

**User Response:** None.

**UDIN1109I: Attempting to save ws.ext.dir changes.**

**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.

**User Response:** None.

**UDIN1110I: Changes saved successfully.**

**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.

**User Response:** None.

**UDIN2001I: Assuming hard coded defaults.**

**Explanation:** This is an informational message issued by the UDDI setup script appserversetupuddi.jacl.

**User Response:** None.

**UDIN2002I: Listing members of type parenttype. Value is:**

**Explanation:** This is an informational message issued by the UDDI setup script appserversetupuddi.jacl.

**User Response:** None.

**UDIN2003I: List for type parenttype returned N members. Values are:**

**Explanation:** This is an informational message issued by the UDDI setup script appserversetupuddi.jacl.

**User Response:** None.

**UDIN2004I: Seeking parenttype with requested id of parentname. Values are:**

**Explanation:** This is an informational message issued by the UDDI setup script appserversetupuddi.jacl.

**User Response:** None.

**UDIN2005I: Checking parentID with parentname. Values are:**

**Explanation:** This is an informational message issued by the UDDI setup script appserversetupuddi.jacl.

**User Response:** None.

**UDIN2006I: Using this as parenttype of parentname. Values are:**

**Explanation:** This is an informational message issued by the UDDI setup script appserversetupuddi.jacl.

**User Response:** None.

**UDIN2007I: Checking for existing childtype under parentname. Values are:**

**Explanation:** This is an informational message issued by the UDDI setup script appserversetupuddi.jacl.

**User Response:** None.

**UDIN2008I: List for childtype returned N members. Values are:**

**Explanation:** This is an informational message issued by the UDDI setup script appserversetupuddi.jacl.

**User Response:** None.

**UDIN2009I: No existing childtype present. Value is:**

**Explanation:** This is an informational message issued by the UDDI setup script appserversetupuddi.jacl.

**User Response:** None.

**UDIN2010I: N existing objects of type childtype found, examining for conflict with childname.**

**Values are:**

**Explanation:** This is an informational message issued by the UDDI setup script appserversetupuddi.jacl.

**User Response:** None.

**UDIN2011I: Checking childID with name childname. Values are:**

**Explanation:** This is an informational message issued by the UDDI setup script appserversetupuddi.jacl.

**User Response:** None.

**UDIN2012I: Conflict found with existing childtype of id childID. Values are:**

**Explanation:** This is an informational message issued by the UDDI setup script appserversetupuddi.jacl.

**User Response:** None.

**UDIN2013I: Not creating requested object of type childtype. Value is:**

**Explanation:** This is an informational message issued by the UDDI setup script appserversetupuddi.jacl.

**User Response:** None.

**UDIN2014I: Conflict found with existing childtype, removing existing childtype. Value is:**

**Explanation:** This is an informational message issued by the UDDI setup script appserversetupuddi.jacl.

**User Response:** None.

**UDIN2015I: Removal of childtype was successful. Value is:**

**Explanation:** This is an informational message issued by the UDDI setup script appserversetupuddi.jacl.

**User Response:** None.

**UDIN2016I: Not in conflict.**

**Explanation:** This is an informational message issued by the UDDI setup script appserversetupuddi.jacl.

**User Response:** None.

**UDIN2017I: Attempting to create childtype under parentname of parentID. Values are:**

**Explanation:** This is an informational message issued by the UDDI setup script appserversetupuddi.jacl.

**User Response:** None.

**UDIN2018I: Create command that will be issued is:**

**Explanation:** This is an informational message issued by the UDDI setup script appserversetupuddi.jacl.

**User Response:** None.

**UDIN2019I: childtype childID was successfully created. Values are:**

**Explanation:** This is an informational message issued by the UDDI setup script appserversetupuddi.jacl.

**User Response:** None.

**UDIN2020I: No matches found.**

**Explanation:** This is an informational message issued by the UDDI setup script appserversetupuddi.jacl.

**User Response:** None.

**UDIN2021I: Looking for builtin\_rra.**

**Explanation:** This is an informational message issued by the UDDI setup script appserversetupuddi.jacl.

**User Response:** None.



**UDIN2022I: List for J2CResourceAdapter returned N members. Value is:**

**Explanation:** This is an informational message issued by the UDDI setup script appserversetupuddi.jacl.

**User Response:** None.

**UDIN2023I: Hunting J2CResourceAdapter associated with Node nodename. Value is:**

**Explanation:** This is an informational message issued by the UDDI setup script appserversetupuddi.jacl.

**User Response:** None.

**UDIN2024I: Using rraID as builtin\_rra. Value is:**

**Explanation:** This is an informational message issued by the UDDI setup script appserversetupuddi.jacl.

**User Response:** None.

**UDIN2025I: Using provider class of implclass with a classpath of clpath. Values are:**

**Explanation:** This is an informational message issued by the UDDI setup script appserversetupuddi.jacl.

**User Response:** None.

**UDIN2026I: Installing to node nodename using database type of dbtype. Values are:**

**Explanation:** This is an informational message issued by the UDDI setup script appserversetupuddi.jacl.

**User Response:** None.

**UDIN2027I: Attempting to create UDDI JDBCProvider.**

**Explanation:** This is an informational message issued by the UDDI setup script appserversetupuddi.jacl.

**User Response:** None.

**UDIN2028I: Attempting to create UDDI Datasource.**

**Explanation:** This is an informational message issued by the UDDI setup script appserversetupuddi.jacl.

**User Response:** None.

**UDIN2029I: Application Manager appmgr found. Value is:**

**Explanation:** This is an informational message issued by the UDDI setup script appserversetupuddi.jacl.

**User Response:** None.

**UDIN2030I: Attempting to install UDDI Registry application.**

**Explanation:** This is an informational message issued by the UDDI setup script appserversetupuddi.jacl.

**User Response:** None.

**UDIN2031I: Checking for installed UDDI Registry application of name appname. Value is:**

**Explanation:** This is an informational message issued by the UDDI setup script appserversetupuddi.jacl.

**User Response:** None.

**UDIN2032I: List for Applications returned N members. Value is:**

**Explanation:** This is an informational message issued by the UDDI setup script appserversetupuddi.jacl.

**User Response:** None.

**UDIN2033W: Application of name appname is not present. Value is:**

**Explanation:** This is an informational message issued by the UDDI setup script appserversetupuddi.jacl.

**User Response:** None.



**UDIN2034I: ApplicationManager not running, so application will not need to be stopped.**

**Explanation:** This is an informational message issued by the UDDI setup script appserversetupuddi.jacl.

**User Response:** None.

**UDIN2035I: Stopping application of name appname. Value is:**

**Explanation:** This is an informational message issued by the UDDI setup script appserversetupuddi.jacl.

**User Response:** None.

**UDIN2036W: stopApplication command for application appname caught exception Exc. Application might not have been running on this server. Values are:**

**Explanation:** This is an informational message issued by the UDDI setup script appserversetupuddi.jacl.

**User Response:** None.

**UDIN2037I: Application appname stopped successfully. Value is:**

**Explanation:** This is an informational message issued by the UDDI setup script appserversetupuddi.jacl.

**User Response:** None.

**UDIN2038I: Removing application appname. Value is:**

**Explanation:** This is an informational message issued by the UDDI setup script appserversetupuddi.jacl.

**User Response:** None.

**UDIN2039I: Application appname removed successfully. Value is:**

**Explanation:** This is an informational message issued by the UDDI setup script appserversetupuddi.jacl.

**User Response:** None.

**UDIN2040I: Attempting to install application appname. Value is:**

**Explanation:** This is an informational message issued by the UDDI setup script appserversetupuddi.jacl.

**User Response:** None.

**UDIN2041I: Starting UDDI application.**

**Explanation:** This is an informational message issued by the UDDI setup script appserversetupuddi.jacl.

**User Response:** None.

**UDIN2042I: Application appname started successfully. Value is:**

**Explanation:** This is an informational message issued by the UDDI setup script appserversetupuddi.jacl.

**User Response:** None.

**UDIN2101I: Attempting to save new configuration.**

**Explanation:** This is an informational message issued by the UDDI setup script appserversetupuddi.jacl.

**User Response:** None.

**UDIN2102I: Changes saved successfully.**

**Explanation:** This is an informational message issued by the UDDI setup script appserversetupuddi.jacl.

**User Response:** None.

**UDIN2103I: Changes were not saved on this call.**

**Explanation:** This is an informational message issued by the UDDI setup script appserversetupuddi.jacl.

**User Response:** None.

**UDIN2104I: Attempting to save post installation configuration.**

**Explanation:** This is an informational message issued by the UDDI setup script appserversetupuddi.jacl.

**User Response:** None.

**UDIN2105I: Changes saved successfully for UDDI Registry.**

**Explanation:** This is an informational message issued by the UDDI setup script appserversetupuddi.jacl.

**User Response:** None.

**UDIN2106I: Attempting to save new configuration.**

**Explanation:** This is an informational message issued by the UDDI setup script appserversetupuddi.jacl.

**User Response:** None.

**UDIN2107I: Changes saved successfully for UDDI Registry.**

**Explanation:** This is an informational message issued by the UDDI setup script appserversetupuddi.jacl.

**User Response:** None.

**UDIN3001I: Application Manager appmgr found. Value is:**

**Explanation:** This is an informational message issued by the UDDI setup script appserverremoveuddi.jacl.

**User Response:** None.

**UDIN3002I: Checking for installed UDDI Registry application.**

**Explanation:** This is an informational message issued by the UDDI setup script appserverremoveuddi.jacl.

**User Response:** None.

**UDIN3003I: List for Applications returned N members. Value is:**

**Explanation:** This is an informational message issued by the UDDI setup script appserverremoveuddi.jacl.

**User Response:** None.

**UDIN3004W: Application of name appname is not present. Value is:**

**Explanation:** This is an informational message issued by the UDDI setup script appserverremoveuddi.jacl.

**User Response:** None.

**UDIN3005I: ApplicationManager not running, so application will not need to be stopped.**

**Explanation:** This is an informational message issued by the UDDI setup script appserverremoveuddi.jacl.

**User Response:** None.

**UDIN3006I: Stopping application of name appname. Value is:**

**Explanation:** This is an informational message issued by the UDDI setup script appserverremoveuddi.jacl.

**User Response:** None.

**UDIN3007W: stopApplication command for application appname caught exception Exc. Application might not have been running on this server. Values are:**

**Explanation:** This is an informational message issued by the UDDI setup script appserverremoveuddi.jacl.

**User Response:** None.

**UDIN3008I: Application appname stopped successfully. Value is:**

**Explanation:** This is an informational message issued by the UDDI setup script appserverremoveuddi.jacl.

**User Response:** None.

**UDIN3009I: Removing application appname. Value is:**

**Explanation:** This is an informational message issued by the UDDI setup script appserverremoveuddi.jacl.

**User Response:** None.

**UDIN3010I: Application appname removed successfully. Value is:**

**Explanation:** This is an informational message issued by the UDDI setup script appserverremoveuddi.jacl.

**User Response:** None.

**UDIN3011I: UDDI datasource will be removed from server servername in node nodename if present. Values are:**

**Explanation:** This is an informational message issued by the UDDI setup script appserverremoveuddi.jacl.

**User Response:** None.

**UDIN3012I: Removing UDDI datasource.**

**Explanation:** This is an informational message issued by the UDDI setup script appserverremoveuddi.jacl.

**User Response:** None.

**UDIN3013I: UDDI datasource successfully removed.**

**Explanation:** This is an informational message issued by the UDDI setup script appserverremoveuddi.jacl.

**User Response:** None.

**UDIN3014I: UDDI JDBC driver will be removed from server servername in node nodename if present. Values are:**

**Explanation:** This is an informational message issued by the UDDI setup script appserverremoveuddi.jacl.

**User Response:** None.

**UDIN3015I: Removing UDDI JDBC driver from node nodename. Value is:**

**Explanation:** This is an informational message issued by the UDDI setup script appserverremoveuddi.jacl.

**User Response:** None.

**UDIN3016I: UDDI JDBC driver successfully removed.**

**Explanation:** This is an informational message issued by the UDDI setup script appserverremoveuddi.jacl.

**User Response:** None.

**UDIN3017I: UDDI Registry application, JDBC driver and datasource removed successfully.**

**Explanation:** This is an informational message issued by the UDDI setup script appserverremoveuddi.jacl.

**User Response:** None.

**UDIN3101I: Attempting to save new configuration.**

**Explanation:** This is an informational message issued by the UDDI setup script appserverremoveuddi.jacl.

**User Response:** None.

**UDIN3102I: Changes to remove UDDI Registry have been saved successfully.**

**Explanation:** This is an informational message issued by the UDDI setup script appserverremoveuddi.jacl.

**User Response:** None.

**UDIN3103I: Attempting to save new configuration.**

**Explanation:** This is an informational message issued by the UDDI setup script appserverremoveuddi.jacl.

**User Response:** None.

**UDIN3104I: Changes saved successfully.**

**Explanation:** This is an informational message issued by the UDDI setup script appserverremoveuddi.jacl.

**User Response:** None.

**UDIN3105I: Attempting final save of new configuration.**

**Explanation:** This is an informational message issued by the UDDI setup script appserverremoveuddi.jacl.

**User Response:** None.

**UDIN3106I: Changes saved successfully.**

**Explanation:** This is an informational message issued by the UDDI setup script appserverremoveuddi.jacl.

**User Response:** None.

**UDIN6001E: This script must be run in a Deployment Manager environment.**

**Explanation:** This is an error message issued by the UDDI setup script setupuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN6002E: To install in a standalone application server, use appserversetupuddi.jacl.**

**Explanation:** This is an error message issued by the UDDI setup script setupuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN6003E: Incorrect number of arguments passed to script.**

**Explanation:** This is an error message issued by the UDDI setup script setupuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN6004E: Usage is:**

**Explanation:** This is an error message issued by the UDDI setup script setupuddi.jacl. The text following 'Usage is:' gives the syntax for calling setupuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN6005E: (<db2userid> <db2password> <db2ziplocation> are only required if setting up to use DB2).** **Explanation:** This is an error message issued by the UDDI setup script setupuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN6006E: Use all forward (/) slashes to avoid problems with escaping back (\) slashes.**

**Explanation:** This is an error message issued by the UDDI setup script setupuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN6007E: Removal of childtype childname caught exception Exc. Values are:**

**Explanation:** This is an error message issued by the UDDI setup script setupuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN6008E: An exception Exc occurred while creating childtype. Values are:**

**Explanation:** This is an error message issued by the UDDI setup script setupuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN6009E: Unable to find requested parenttype of parentname. Values are:**

**Explanation:** This is an error message issued by the UDDI setup script setupuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN6010E: List command for J2CResourceAdapter caught exception Exc. Value is:**

**Explanation:** This is an error message issued by the UDDI setup script setupuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN6011E: No J2CResourceAdapter objects available.**

**Explanation:** This is an error message issued by the UDDI setup script setupuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN6012E: An error occurred during execution of setupuddi.jacl. Please check the parameters and try again.**

**Explanation:** This is an error message issued by the UDDI setup script setupuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN6013E: Uninstall of application appname caught exception Exc. Values are:**  
**Explanation:** This is an error message issued by the UDDI setup script setupuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN6014E: Install of UDDI application caught exception Exc. Value is:**  
**Explanation:** This is an error message issued by the UDDI setup script setupuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN6015E: Could not get JVM.**  
**Explanation:** This is an error message issued by the UDDI setup script setupuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN6016E: Cannot find nodeSync MBean.**  
**Explanation:** This is an error message issued by the UDDI setup script setupuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN6017E: nodeSync failed. UDDI Application may not be fully installed.**  
**Explanation:** This is an error message issued by the UDDI setup script setupuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN6018E: stopServer command for server servername caught exception Exc. Values are:**  
**Explanation:** This is an error message issued by the UDDI setup script setupuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN6019E: startServer command for server servername caught exception Exc. Values are:**  
**Explanation:** This is an error message issued by the UDDI setup script setupuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN6101E: Error saving configuration, changes not saved due to exception Exc. Value is:**  
**Explanation:** This is an error message issued by the UDDI setup script setupuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN6102E: Error saving configuration, changes not saved due to exception Exc. Value is:**  
**Explanation:** This is an error message issued by the UDDI setup script setupuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN6103E: Error saving configuration, changes not saved due to exception Exc. Value is:**  
**Explanation:** This is an error message issued by the UDDI setup script setupuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN6104E: Error saving configuration, changes not saved due to exception Exc. Value is:**  
**Explanation:** This is an error message issued by the UDDI setup script setupuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN7001E: This script must be run in a Deployment Manager environment.**  
**Explanation:** This is an error message issued by the UDDI setup script removeuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN7002E: To remove from a standalone application server, use appserverremoveuddi.jacl.**  
**Explanation:** This is an error message issued by the UDDI setup script removeuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN7003E: Incorrect number of arguments passed to script.**  
**Explanation:** This is an error message issued by the UDDI setup script removeuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN7004E: Usage is:**  
**Explanation:** This is an error message issued by the UDDI setup script removeuddi.jacl. The text following 'Usage is:' gives the syntax for calling removeuddi.jacl.

**User Response:**This message is self-explanatory.

- UDIN7005E: stopServer command for server servername caught exception Exc. Values are:**  
**Explanation:** This is an error message issued by the UDDI setup script removeuddi.jacl.  
**User Response:**This message is self-explanatory.
- UDIN7006E: Removal of resource bundles caught exception Exc. Value is:**  
**Explanation:** This is an error message issued by the UDDI setup script removeuddi.jacl.  
**User Response:**This message is self-explanatory.
- UDIN7007E: Removal of Cloudscape user functions caught exception Exc. Value is:**  
**Explanation:** This is an error message issued by the UDDI setup script removeuddi.jacl.  
**User Response:**This message is self-explanatory.
- UDIN7008E: Uninstall of application appname caught exception Exc. Values are:**  
**Explanation:** This is an error message issued by the UDDI setup script removeuddi.jacl.  
**User Response:**This message is self-explanatory.
- UDIN7009E: Removal of UDDI datasource caught exception Exc. Value is:**  
**Explanation:** This is an error message issued by the UDDI setup script removeuddi.jacl.  
**User Response:**This message is self-explanatory.
- UDIN7010E: Removal of UDDI JDBC driver caught exception Exc. Value is:**  
**Explanation:** This is an error message issued by the UDDI setup script removeuddi.jacl.  
**User Response:**This message is self-explanatory.
- UDIN7011E: Removal of configuration properties file caught exception Exc. Value is:**  
**Explanation:** This is an error message issued by the UDDI setup script removeuddi.jacl.  
**User Response:**This message is self-explanatory.
- UDIN7012E: Cannot find nodeSync MBean.**  
**Explanation:** This is an error message issued by the UDDI setup script removeuddi.jacl.  
**User Response:**This message is self-explanatory.
- UDIN7013E: nodeSync failed. UDDI Application may not be fully uninstalled.**  
**Explanation:** This is an error message issued by the UDDI setup script removeuddi.jacl.  
**User Response:**This message is self-explanatory.
- UDIN7014E: startServer command for server servername caught exception Exc. Values are:**  
**Explanation:** This is an error message issued by the UDDI setup script removeuddi.jacl.  
**User Response:**This message is self-explanatory.
- UDIN7015E: Could not get JVM.**  
**Explanation:** This is an error message issued by the UDDI setup script removeuddi.jacl.  
**User Response:**This message is self-explanatory.
- UDIN7101E: Error saving configuration, changes not saved due to exception Exc. Value is:**  
**Explanation:** This is an error message issued by the UDDI setup script removeuddi.jacl.  
**User Response:**This message is self-explanatory.
- UDIN7102E: Error saving configuration, changes not saved due to exception Exc. Value is:**  
**Explanation:** This is an error message issued by the UDDI setup script removeuddi.jacl.  
**User Response:**This message is self-explanatory.
- UDIN7103E: Error saving configuration, changes not saved due to exception Exc. Value is:**  
**Explanation:** This is an error message issued by the UDDI setup script removeuddi.jacl.  
**User Response:**This message is self-explanatory.
- UDIN7104E: Error saving configuration, changes not saved due to exception Exc. Value is:**  
**Explanation:** This is an error message issued by the UDDI setup script removeuddi.jacl.  
**User Response:**This message is self-explanatory.
- UDIN7105E: Error saving configuration, changes not saved due to exception Exc. Value is:**  
**Explanation:** This is an error message issued by the UDDI setup script removeuddi.jacl.  
**User Response:**This message is self-explanatory.



- UDIN8001E: This script must be run on a standalone application server.**  
**Explanation:** This is an error message issued by the UDDI setup script appserversetupuddi.jacl.  
**User Response:**This message is self-explanatory.
- UDIN8002E: To install in a Deployment Manager Environment, use setupuddi.jacl.**  
**Explanation:** This is an error message issued by the UDDI setup script appserversetupuddi.jacl.  
**User Response:**This message is self-explanatory.
- UDIN8003E: Incorrect number of arguments passed to script.**  
**Explanation:** This is an error message issued by the UDDI setup script appserversetupuddi.jacl.  
**User Response:**This message is self-explanatory.
- UDIN8004E: Usage is:**  
**Explanation:** This is an error message issued by the UDDI setup script appserversetupuddi.jacl.  
The text following 'Usage is:' gives the syntax for calling appserversetupuddi.jacl.  
**User Response:**This message is self-explanatory.
- UDIN8005E: (<db2userid> <db2password> <db2ziplocation> are only required if setting up to use DB2).** **Explanation:** This is an error message issued by the UDDI setup script appserversetupuddi.jacl.  
**User Response:**This message is self-explanatory.
- UDIN8006E: Use all forward (/) slashes to avoid problems with escaping back (\) slashes.**  
**Explanation:** This is an error message issued by the UDDI setup script appserversetupuddi.jacl.  
**User Response:**This message is self-explanatory.
- UDIN8007E: List command for type parenttype caught exception Exc. Values are:**  
**Explanation:** This is an error message issued by the UDDI setup script appserversetupuddi.jacl.  
**User Response:**This message is self-explanatory.
- UDIN8008E: No objects of type parenttype available. Value is:**  
**Explanation:** This is an error message issued by the UDDI setup script appserversetupuddi.jacl.  
**User Response:**This message is self-explanatory.
- UDIN8009E: List command for childtype caught exception Exc. Values are:**  
**Explanation:** This is an error message issued by the UDDI setup script appserversetupuddi.jacl.  
**User Response:**This message is self-explanatory.
- UDIN8001E: This script must be run on a standalone application server.**  
**Explanation:** This is an error message issued by the UDDI setup script appserversetupuddi.jacl.  
**User Response:**This message is self-explanatory.
- UDIN8010E: Error during remove of existing childtype, exception Exc caught. Values are:**  
**Explanation:** This is an error message issued by the UDDI setup script appserversetupuddi.jacl.  
**User Response:**This message is self-explanatory.
- UDIN8011E: Create command for childtype caught exception Exc. Values are:**  
**Explanation:** This is an error message issued by the UDDI setup script appserversetupuddi.jacl.  
**User Response:**This message is self-explanatory.
- UDIN8012E: Unable to find requested parenttype of parentname. Values are:**  
**Explanation:** This is an error message issued by the UDDI setup script appserversetupuddi.jacl.  
**User Response:**This message is self-explanatory.
- UDIN8013E: List command for J2CResourceAdapter caught exception Exc. Value is:**  
**Explanation:** This is an error message issued by the UDDI setup script appserversetupuddi.jacl.  
**User Response:**This message is self-explanatory.
- UDIN8014E: No J2CResourceAdapter objects available.**  
**Explanation:** This is an error message issued by the UDDI setup script appserversetupuddi.jacl.  
**User Response:**This message is self-explanatory.



**UDIN8015E: An error occurred during execution of appserversetupuddi.jacl. Please check the parameters and try again.**

**Explanation:** This is an error message issued by the UDDI setup script appserversetupuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN8016E: List command for Applications caught exception Exc. Value is:**

**Explanation:** This is an error message issued by the UDDI setup script appserversetupuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN8017E: Uninstall of application appname caught exception Exc. Values are:**

**Explanation:** This is an error message issued by the UDDI setup script appserversetupuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN8018E: Install of UDDI application caught exception Exc. Value is:**

**Explanation:** This is an error message issued by the UDDI setup script appserversetupuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN8019E: startApplication command for appname caught exception Exc. Values are:**

**Explanation:** This is an error message issued by the UDDI setup script appserversetupuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN8101E: Error saving configuration, changes not saved due to exception Exc. Value is:**

**Explanation:** This is an error message issued by the UDDI setup script appserversetupuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN8102E: Error saving configuration, changes not saved due to exception Exc. Value is:**

**Explanation:** This is an error message issued by the UDDI setup script appserversetupuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN8103E: Error saving configuration, changes not saved due to exception Exc. Value is:**

**Explanation:** This is an error message issued by the UDDI setup script appserversetupuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN9001E: This script must be run on a standalone application server.**

**Explanation:** This is an error message issued by the UDDI setup script appserverremoveuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN9002E: To remove from a deployment manager environment, use removeuddi.jacl.**

**Explanation:** This is an error message issued by the UDDI setup script appserverremoveuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN9003E: Incorrect number of arguments passed to script.**

**Explanation:** This is an error message issued by the UDDI setup script appserverremoveuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN9004E: Usage is:**

**Explanation:** This is an error message issued by the UDDI setup script appserverremoveuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN9005E: List command for Applications caught exception Exc. Value is:**

**Explanation:** This is an error message issued by the UDDI setup script appserverremoveuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN9006E: Uninstall of application appname caught exception Exc. Values are:**

**Explanation:** This is an error message issued by the UDDI setup script appserverremoveuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN9007E: Removal of UDDI datasource caught exception Exc. Value is:**

**Explanation:** This is an error message issued by the UDDI setup script appserverremoveuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN9008E: Removal of UDDI JDBC driver caught exception Exc. Value is:**

**Explanation:** This is an error message issued by the UDDI setup script appserverremoveuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN9101E: Error saving configuration, changes not saved due to exception Exc. Value is:**

**Explanation:** This is an error message issued by the UDDI setup script appserverremoveuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN9102E: Error saving configuration, changes not saved due to exception Exc. Value is:**

**Explanation:** This is an error message issued by the UDDI setup script appserverremoveuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN9103E: Error saving configuration, changes not saved due to exception Exc. Value is:**

**Explanation:** This is an error message issued by the UDDI setup script appserverremoveuddi.jacl.

**User Response:**This message is self-explanatory.

## UDLC (Web Services UDDI) messages

There are no messages issued by this component.

## UDPR (Web Services UDDI) messages

There are no messages issued by this component.

## UDRS (Web Services UDDI) messages

**UDRS0001E: Exception "<exception>" occurred while attempting to get UDDI Message Logger.**

**Explanation:** This message is issued to stderr when an attempt to get the UDDI Message Logger fails with the indicated exception. Since the attempt to get the message logger failed, the message cannot be logged. No messages can be logged by this instance of the IBM WebSphere UDDI Registry.

**User Response:** Restart the UDDI registry. If the error persists, examine the WebSphere logs for information on its cause. If the problem cannot be resolved, then please contact the IBM Customer Service Center.

**UDRS0002E: Exception "<exception>" occurred while attempting to get UDDI Trace Logger for "<component>".**

**Explanation:** This message is logged when an attempt to get the UDDI Trace Logger for the specified component (or package) fails with the indicated exception. No trace entries can be logged for this component or package of the IBM WebSphere UDDI Registry.

**User Response:** Restart the UDDI registry. If the error persists, examine the WebSphere logs for information on its cause. If the problem cannot be resolved, then please contact the IBM Customer Service Center.

## UDSC (Web Services UDDI) messages

There are no messages issued by this component.

## UDSP (Web Services UDDI) messages

**UDSP0001E: ParserPool found empty whilst attempting to process request. Request unsatisfied**

**Explanation:** A SOAP request was received, but was unable to be dealt with, as there were no free Parsers within the ParserPool.

**User Response:** Consider increasing the number of Parsers within the ParserPool by modifying the Init Parameter on the SOAP servlets.

**UDSP0002E: Error locating schemas required for UDDI processing. SOAP Servlets unworkable.**

**Explanation:** The SOAP servlet was unable to locate the schemas it requires in order to process SOAP requests. Without these, the servlet cannot process SOAP requests.

**User Response:** Check installation of UDDI was performed correctly. If the error persists, examine the WebSphere logs for information on its cause. If the problem cannot be resolved, then please contact the IBM Customer Service Center.

**UDSP0003W: Servlet unable to locate init parameter 'defaultPoolSize'. Using internal defaults.**

**Explanation:** The SOAP servlet was unable to locate the init parameter which sets the default size of the ParserPool. It will fall back to an internal default.

**User Response:** If this message occurred after attempting to make changes to the defaultPoolSize init parameter, ensure the changes were correct. If this message has appeared after installed, ensure installation was performed correctly.

**UDSP0004W: Servlet unable to understand init parameter 'defaultPoolSize'. Using internal defaults.**

**Explanation:** The SOAP servlet was unable to parse the init parameter which sets the default size of the ParserPool. It will fall back to an internal default.

**User Response:** If this message occurred after attempting to make changes to the defaultPoolSize init parameter, ensure the changes were correct. If this message has appeared after installed, ensure installation was performed correctly.

**UDSP0005E: Error occurred during parser creation.**

**Explanation:** An unspecified error occurred during the creation of a SOAP parser

**User Response:** Restart the UDDI registry. If the error persists, examine the WebSphere logs for information on its cause. If the problem cannot be resolved, then please contact the IBM Customer Service Center.

**UDSP0006E: Internal configuration error.**

**Explanation:** This error may occur if there was a failure creating a Parser, with accompanying message UDSP0005. It may also occur if there was a problem acquiring the Persistence layer.

**User Response:** Restart the UDDI registry. If the error persists, examine the WebSphere logs for information on its cause. If the problem cannot be resolved, then please contact the IBM Customer Service Center.

**UDSP0007E: Error during servlet acquisition of persistence layer.**

**Explanation:** The SOAP servlet was unable to acquire the persistence layer required for it to communicate with the UDDI datasource

**User Response:** Restart the UDDI registry. If the error persists, examine the WebSphere logs for information on its cause. If the problem cannot be resolved, then please contact the IBM Customer Service Center.

**UDSP0008E: Error during servlet release of persistence layer.**

**Explanation:** The persistence layer reported a problem when the SOAP servlet attempted to release it.

**User Response:** Restart the UDDI registry. If the error persists, examine the WebSphere logs for information on its cause. If the problem cannot be resolved, then please contact the IBM Customer Service Center.

**UDSP0009E: Error during sending of response to client.**

**Explanation:** An error occurred when sending a SOAP response message back to a client. The client may not have received the response

**User Response:** This error is recorded to enable logging of failed responses to clients. The error may be the fault of the client disconnecting before the reply could be sent, or may indicate a network problem. Examine the WebSphere logs for more information on its cause.

## **UDUC (Web Services UDDI) messages**

**UDUC0001I: IBM WebSphere UDDI Registry user console starting initialization.**

**Explanation:** The user console control servlet is starting.

**User Response:** None.

**UDUC0002I: IBM WebSphere UDDI Registry user console finished initialization.**

**Explanation:** The user console control servlet has completed startup.

**User Response:** None.

**UDUC0003I: Reading init parameters.**

**Explanation:** The user console control servlet has started reading external parameters in its init method

**User Response:** None.

**UDUC0004I: Finished reading init parameters.**

**Explanation:** The user console control servlet has finished reading external parameters in its init method. This message indicates the user console is ready to accept client requests.

**User Response:** None.

**UDUC0005E: A serious error has occurred. Error message: <Message> error: <Throwable>. More information: <Additional information>.**

**Explanation:** This error message indicates an unexpected error has occurred. The <Message> describes the error that has occurred and the <Throwable> is the type of error that was caught. <Additional information> may provide further information, if available.

**User Response:** A trace of the gui component is recommended. Contact IBM support with this information.

**UDUC0006E: A persistence error has occurred. Error message: <Message> error: <Throwable>. More information: <Additional information>.**

**Explanation:** An error occurred while performing a database operation. The <Message> describes the error that occurred and the <Throwable> is the type of error that was caught. <Additional information> may provide further information, if available.

**User Response:** Check database connections and state. Please provide IBM support with a trace, including the gui and persistence components.

**UDUC0007E: A User mismatch error has occurred. Error message: <Message> error: <Throwable>. More information: <Additional information>.**

**Explanation:** The user id provided does not match the user id required or expected whilst performing an operation that requires authentication. The <Message> describes the error that occurred and the <Throwable> is the type of error that was caught. <Additional information> may provide further information, if available.

**User Response:** Check the user has authority for the operation being requested. If necessary, contact IBM support detailing the actions taken to recreate the problem.

**UDUC0008E: An invalid key was passed. Error message: <Message> error: <Throwable>. More information: <Additional information>.**

**Explanation:** The requested operation is trying to retrieve information about an entity with a key that is invalid. This may occur if the entity has been deleted by another session. The <Message> describes the error that occurred and the <Throwable> is the type of error that was caught. <Additional information> may provide further information, if available.

**User Response:** Ask the client to close existing sessions and attempt the operation in a new browser session. If the problem persists, please provide IBM support with a trace of the gui and api components.

**UDUC0009E: An invalid value was supplied. Error message: <Message> error: <Throwable>. More information: <Additional information>**

**Explanation:** An invalid value was passed to an API call. The >Message> describes the error that occurred and the <Throwable> is the type of error that was caught. <Additional information> may provide further information, if available.

**User Response:** Contact IBM support with a trace of the gui and api components.

**UDUC0010E: Failed to introspect ActionForm properties. Exception: <Exception>.**

**Explanation:** String properties of a form object could not be introspected which means that the form contents cannot be checked for invalid characters.

**User Response:** Please contact IBM support with details of the Exception and a trace of the gui component.

**UDUC0011E: Failed to invoke reflected methods in ActionForm. Exception: <Exception>.**

**Explanation:** A form object's declared public method for setting or getting a String value could not be invoked. This method is required to check for invalid characters.

**User Response:** Please contact IBM support with details of the Exception and a trace of the gui component.

**UDUC0012E: User console initialization failed to connect to UDDI database. Exception: <Exception>.**

**Explanation:** During user console initialization, connection to the database failed, and threw the exception specified.

**User Response:** Check the connection to the UDDI database. The included exception message may yield some clues to help you resolve the problem. If unresolved, please contact IBM support with a trace of the gui component during startup.

**UDUC0013E: User console initialization failed to initialize tModels. Exception: <Exception>.**

**Explanation:** Indicates that an error has occurred during initialization of ActionServlet, specifically when reading tModels (invoking init method in class TModelNames).

**User Response:** Check the state of the UDDI database. Visually inspect the TMODEL table and confirm it is populated with valid data. The included exception message may yield some clues to help you resolve the problem. If unresolved, please contact IBM support with a trace of the gui component during startup.

**UDUC0014E: User console initialization failed to initialize taxonomies. Exception: <Exception>.**

**Explanation:** Indicates that an error has occurred during initialization of ActionServlet, specifically when reading taxonomy data (invoking init method of CategoryTaxonomyTree).

**User Response:** Check the state of the UDDI database. The included exception message may yield some clues to help you resolve the problem. If unresolved, please contact IBM support with a trace of the gui component during startup.

## UDUT UDDI Utility Tools messages

**UDUT0001I: Usage: java -jar UDDIUtilityTools.jar {'function'} [options]**

**function:**

<code>-promote <i>entity source</i></code>	Promote entities between registries
<code>-export <i>entity source</i></code>	Extract entities from registry to XML
<code>-delete <i>entity source</i></code>	Delete entities from registry
<code>-import</code>	Create entities from XML to registry

**where *entity source* is one of:**

<code>-tmodel</code>	Specify single entity type and key
<code>-business</code>	Specify single entity type and key
<code>-service</code>	Specify single entity type and key
<code>-binding <i>key</i></code>	Specify single entity type and key
<code>-keysFile   -f <i>filename</i></code>	Specify file containing entity types and keys

**options:**

<code>-properties <i>filename</i></code>	Specify path to configuration file
<code>-overwrite   -o</code>	Overwrite an entity if it already exists
<code>-log   -v</code>	Output verbose messages
<code>-definitionFile <i>filename</i></code>	Specify path to UDDI entity definition file
<code>-importReferenced</code>	Import entities referenced by source entities

**The following options override property settings in configuration file:**

<code>-overwrite</code>
<code>-log</code>
<code>-definitionFile</code>
<code>-importReferenced</code>

**Example:** `java -jar UDDIUtilityTools.jar -promote -keysFile C:/uddikeys.txt`

**Explanation:** This is the usage message displayed at the command line when the user has entered an invalid combination of arguments or options.

**User Response:** Enter the command according to the usage message..

**UDUT0002I: \*\*\*\*\* Starting UDDI Utility Tools \*\*\*\*\***  
**Explanation:** This message is used as a marker in the message log file to indicate tool start points.

**User Response:** None.

**UDUT0003I: Promoting entityType<entity type> key<entity key>...**  
**Explanation:** Indicates which entity type (business, tModel and so on) is being promoted, and it's key value.

**User Response:** None.

**UDUT0004I: Bad entityType: received<incorrect entity type>, expected <tModel|business|service|binding>**  
**Explanation:** The user entered an incorrect entity type.

**User Response:** Use an entity type of tModel, business, service or binding.

**UDUT0005I: Promotion successful.**  
**Explanation:** Indicates the promote function completed successfully.

**User Response:** None.

**UDUT0006I: Import successful.**  
**Explanation:** Indicates the import function completed successfully.

**User Response:** None.

**UDUT0007I: Export successful.**  
**Explanation:** Indicates the export function completed successfully.

**User Response:** None.

**UDUT0008I: Delete successful.**  
**Explanation:** Indicates the delete function completed successfully.

**User Response:** None.

**UDUT0009I: Exporting entities ...**  
**Explanation:** Indicates the export function has started.

**User Response:** None.

**UDUT0010I: Exporting business, businessKey[<business key>].**  
**Explanation:** Indicates that the businessEntity with the specified key is being exported.

**User Response:** None.

**UDUT0011I: Exporting service, serviceKey[<service key>].**  
**Explanation:** Indicates that the businessService with the specified key is being exported.

**User Response:** None.

**UDUT0012I: Exporting binding, bindingKey[<binding key>].**  
**Explanation:** Indicates that the bindingTemplate with the specified key is being exported.

**User Response:** None.

**UDUT0013I: Exporting tModel, tModelKey[<tModel key>].**  
**Explanation:** Indicates that the tModel with the specified key is being exported.

**User Response:** None.

**UDUT0014I: Exporting referenced tModel, tModelKey[<tModel key>].**  
**Explanation:** Indicates that the referenced tModel with the specified key is being exported.

**User Response:** None.

**UDUT0015I: Exported <entity count> entities.**  
**Explanation:** Indicates that the export function completed, and shows the number of entities exported.

**User Response:** None.

**UDUT0016I: Importing entities ...**  
**Explanation:** Indicates the import function has started.



- User Response:** None.
- UDUT0017I: Importing business, businessKey[<business key>].**  
**Explanation:** Indicates that the businessEntity with the specified key is being imported.
- User Response:** None.
- UDUT0018I: Importing service, serviceKey[<service key>].**  
**Explanation:** Indicates that the businessService with the specified key is being imported.
- User Response:** None.
- UDUT0019I: Importing binding, bindingKey[<binding key>]**  
**Explanation:** Indicates that the bindingTemplate with the specified key is being imported.
- User Response:** None.
- UDUT0020I: Importing tModel, tModelKey[<tModel key>].**  
**Explanation:** Indicates that the tModel with the specified key is being imported..
- User Response:** None.
- UDUT0021I: Importing referenced tModel, tModelKey[<tModel key>].**  
**Explanation:** Indicates that the referenced tModel with the specified key is being imported.
- User Response:** None.
- UDUT0022I: Imported <entity count> entities.**  
**Explanation:** Indicates that the import function completed, and shows the number of entities imported.
- User Response:** None.
- UDUT0023I: Deleting entities ...**  
**Explanation:** Indicates the delete function has started.
- User Response:** None.
- UDUT0024I: Deleting business, businessKey[<business key>].**  
**Explanation:** Indicates that the businessEntity with the specified key is being deleted.
- User Response:** None.
- UDUT0025I: Deleting service, serviceKey[<service key>].**  
**Explanation:** Indicates that the businessService with the specified key is being deleted.
- User Response:** None.
- UDUT0026I: Deleting binding, bindingKey[<binding key>].**  
**Explanation:** Indicates that the bindingTemplate with the specified key is being deleted.
- User Response:** None.
- UDUT0027I: Deleting tModel, tModelKey[<tModel key>].**  
**Explanation:** Indicates that the tModel with the specified key is being deleted.
- User Response:** None.
- UDUT0028I: Deleted <entity count> entities.**  
**Explanation:** Indicates that the delete function completed, and shows the number of entities deleted.
- User Response:** None.
- UDUT0029I: Serializing ...**  
**Explanation:** Indicates that generation of the Entity Definition File has started.
- User Response:** None.
- UDUT0030I: Serialized entities.**  
**Explanation:** Indicates that generation of the Entity Definition File completed successfully.
- User Response:** None.
- UDUT0031I: Deserializing ...**  
**Explanation:** Indicates that reading of the Entity Definition File and creation of UDDI entities has started.



- User Response:** None.
- UDUT0032I: Deserialized entities.**  
**Explanation:** Indicates that reading of the Entity Definition File and creation of UDDI entities completed successfully.
- User Response:** None.
- UDUT0033I: Function '<function>' completed successfully.**  
**Explanation:** Indicates the requested function completed successfully.
- User Response:** None.
- UDUT0034W: Function '<function>' did not complete successfully. See message log for further information.**  
**Explanation:** Indicates the requested function did not complete successfully.
- User Response:** The message log may yield further information if the verbose option is on. Check the configuration properties file setting are correct. If that does not identify the problem, try running with trace logging enabled. If that does not yield a solution, contact your IBM support center.
- UDUT0035W: Parser error: <warning description>**  
**Explanation:** The XML parser reports a warning about the content of the Entity Definition File.
- User Response:** Based on the context of the warning message, check the validity of the Entity Definition File.
- UDUT0036E: Parser error <error description>**  
**Explanation:** The XML parser reports an error about the content of the Entity Definition File.
- User Response:** Based on the context of the error message, check the validity of the Entity Definition File.
- UDUT0037E: Unrecognized parser feature: <feature description>**  
**Explanation:** A parser feature set by the UDDI Utility Tools is not recognized by the parser.
- User Response:** Check you are using the correct type and level of XML parser. If correct, contact your IBM support center.
- UDUT0038E: Unsupported parser feature: <feature description>**  
**Explanation:** A parser feature set by the UDDI Utility Tools is not supported by the parser.
- User Response:** Check you are using the correct type and level of XML parser. If correct, contact your IBM support center.
- UDUT0039E: Unrecognized parser property: <property description>, value: <value>**  
**Explanation:** A parser property set by the UDDI Utility Tools is not recognized by the parser.
- User Response:** Check you are using the correct type and level of XML parser. If correct, contact your IBM support center.
- UDUT0040E: Unsupported parser property: <property description>, value: <value>**  
**Explanation:** A parser property set by the UDDI Utility Tools is not supported by the parser.
- User Response:** Check you are using the correct type and level of XML parser. If correct, contact your IBM support center.
- UDUT0041I: <message>**  
**Explanation:** This is a placeholder message used during development only.
- User Response:** None.
- UDUT0042E: Unable to find the configuration file: <filepath>**  
**Explanation:** UDDI Utility Tools cannot locate the specified configuration file.
- User Response:** UDDI Utility Tools looks for a default configuration properties with the file name 'UDDIUtilityTools.properties' in the current directory. Check that the configuration file has this name, or that the argument value supplied with the '-properties' option is pointing at a file that exists.
- UDUT0043E: An Exception occurred trying to read the configuration file.**  
**Explanation:** The configuration file could not be read.

**User Response:** Check the file path points to a valid file and that the file does not have the 'hidden' attribute set.

**UDUT0044W: Configuration file is missing the '<property name>' property.**

**Explanation:** A required property is missing from the configuration file.

**User Response:** Add the missing property name and value to the configuration file. Check that the property name is not misspelled.

**UDUT0045W: Property: '<property name>' has value '<property value>'. It must be either 'true' or 'false'.**

**Explanation:** A value was given to a property other than 'true' or 'false'.

**User Response:** Set the property value to 'true' or 'false'.

**UDUT0046W: Property: '<property name>' has value '<property value>'. It must be an integer value.** **Explanation:** A value was given to a property other than an integer value.

**User Response:** Set the property value to an integer value.

**UDUT0047E: Unable to find the keyFile file: <keys file path>**

**Explanation:** The keys file could not be located at the specified path.

**User Response:** Check the file name and path and correct and that the file exists.

**UDUT0048E: Unable to read the keyFile file: <keys file path>**

**Explanation:** The keys file could not be read due to an IO error.

**User Response:** Check the file's hidden attribute is not set.

**UDUT0049E: Unable to write to entity definition file: <entity definition file path>**

**Explanation:** During initialization, the Entity Definition File could not be written to.acl.

**User Response:** Check the file's read only attribute is not set.

**UDUT0050E: Unable to find UDDI Entity definition file: <entity definition file path>**

**Explanation:** The Entity Definition File could not be found at the specified file path.

**User Response:** Check the file path is correct and that the file exists.

**UDUT0051E: Unable to read UDDI Entity definition file: <entity definition file path>**

**Explanation:** The Entity Definition File could not be read due to an IO error.

**User Response:** Check the file's hidden attribute is not set.

**UDUT0052E: Unable to close the message file: <file path>**

**Explanation:** The attempt to close the message file failed.

**User Response:** The disk might be full. If so, clear some space or direct log output to a different disk.

**UDUT0053E: Unable to close the trace file: <file path>**

**Explanation:** The attempt to close the trace log file failed.

**User Response:** The disk might be full. If so, clear some space or direct log output to a different disk.

**UDUT0054E: The logger was unable to find the file: <file path>**

**Explanation:** The UDDI Utility Tools logger could not find the specified file.

**User Response:** None.

**UDUT0055E: ERROR OCCURRED ...**

**Explanation:** General purpose error message used in development only.

**User Response:** None.

**UDUT0056E: Exception:**

**Explanation:** General purpose message prefix used for reporting exceptions.

**User Response:** None.

**UDUT0057W: Only one function may be specified on the command line.**

**Explanation:** Multiple function commands were entered on the command line.

**User Response:** Specify one function in accordance with the usage message.

**UDUT0058W: No function was specified.**

**Explanation:** UDDI Utility Tools was invoked with no function specified.

**User Response:** Specify one function in accordance with the usage message.

**UDUT0059W: The function: <function> was not recognized.**

**Explanation:** The function value did not match any of the allowed functions.

**User Response:** Specify one function in accordance with the usage message.

**UDUT0060W: The argument '<argument>' was not recognized.**

**Explanation:** The argument value does not match any of the allowed arguments.

**User Response:** Specify arguments in accordance with the usage message.

**UDUT0061W: There was a missing value for <argument> argument.**

**Explanation:** An expected value for the specified argument was not supplied.

**User Response:** Specify a value for the argument in accordance with the usage message.

**UDUT0062W: Unexpected argument: <argument> (entity key file is already specified).**

**Explanation:** The entity type argument cannot be specified if the keysFile argument is already specified.

**User Response:** Specify arguments in accordance with the usage message.

**UDUT0063W: Unexpected argument: <argument> (entity key is already specified).**

**Explanation:** The keysFile argument cannot be specified if an entity type argument and key value has already been specified.

**User Response:** Specify arguments in accordance with the usage message.

**UDUT0064W: Argument: <argument> cannot be specified more than once.**

**Explanation:** An argument was specified twice in the same command.

**User Response:** Specify arguments in accordance with the usage message.

**UDUT0065E: No entity keys were specified.**

**Explanation:** A keys file or an entity type and key value must be specified for functions using keys.

**User Response:** Specify arguments in accordance with the usage message.

**UDUT0066E: Could not load Database driver: dbDriver<database driver>.**

**Explanation:** The specified database driver could not be loaded.

**User Response:** Check the database driver value in the configuration file is valid, and the driver's class is present in the classpath property.

**UDUT0067E: Could not create Database connection: dbUrl, dbUser, (dbPasswd not shown).**

**Explanation:** A connection could not be established with the database at the specified URL with the specified userid.

**User Response:** Check the database URL, userid and password values are correct in the configuration file, and that the database manager is running.

**UDUT0068E: Could not close the database connection.**

**Explanation:** An attempt to close the database connection failed.

**User Response:** If the problem persists, contact your IBM support center.

**UDUT0069E: Could not create minimal entity for tModel.**

**Explanation:** The minimal data necessary for a valid tModel could not be inserted in the target UDDI registry database.

**User Response:** Check the database URL, userid and password values are correct in the configuration file, and that the database manager is running.

**UDUT0070E: Could not create minimal entity for Service.**

**Explanation:** The minimal data necessary for a valid businessService could not be inserted in the target UDDI registry database.

**User Response:** Check the database URL, userid and password values are correct in the configuration file, and that the database manager is running.

**UDUT0071E: Could not create minimal entity for Business.**

**Explanation:** The minimal data necessary for a valid businessEntity could not be inserted in the target UDDI registry database.

**User Response:** Check the database URL, userid and password values are correct in the configuration file, and that the database manager is running.

**UDUT0072E: Could not create minimal entity for Binding.**

**Explanation:** The minimal data necessary for a valid bindingTemplate could not be inserted in the target UDDI registry database.

**User Response:** Check the database URL, userid and password values are correct in the configuration file, and that the database manager is running.

**UDUT0073E: There was an error while trying to create an XML Document.**

**Explanation:** An attempt to create the Entity Definition File failed.

**User Response:** Check the file path specified in the configuration file for the Entity Definition File is valid and is not set to read only.

**UDUT0074E: There was an error parsing the entity definition file.**

**Explanation:** An unspecified error occurred when parsing the Entity Definition File.

**User Response:** Check the entity definition file content is valid according to the UDDI Utility Tools schema file, promoter.xsd.

**UDUT0075E: One or more errors occurred while parsing the entity definition file. See message log for details.**

**Explanation:** Errors occurred when parsing the Entity Definition File.

**User Response:** Check the entity definition file content is valid according to the UDDI Utility Tools schema file, promoter.xsd.

**UDUT0076W: One or more warnings were raised while parsing the entity definition file. See message log for details.**

**Explanation:** Warnings occurred when parsing the Entity Definition File.

**User Response:** Check the entity definition file content is valid according to the UDDI Utility Tools schema file, promoter.xsd.

**UDUT0078E: Unable to obtain authinfo.**

**Explanation:** AuthInfo could not be obtained from the UDDI registry with the given userid and password.

**User Response:** Check the userid and password property values are correct in the configuration file.

**UDUT0079E: The inquiryURL is malformed: <inquiry URL>.**

**Explanation:** The inquiry URL specified in the configuration file is not valid.

**User Response:** Correct the value for the inquiry URLs (fromInquiryURL and toInquiryURL) in the configuration file.

**UDUT0080E: The publishURL is malformed: <publish URL>**

**Explanation:** The publish URL specified in the configuration file is not valid.

**User Response:** Correct the value for the publish URL (toPublishURL) in the configuration file.

**UDUT0081E: Could not get tModel detail for tModelKey[<tModel key>].**

**Explanation:** The get tModel operation failed on the source registry.

**User Response:** Check the key exists in the source registry.

**UDUT0082E: Could not get service detail for serviceKey[<service key>].**

**Explanation:** The get service operation failed on the source registry.

**User Response:** Check the key exists in the source registry.

**UDUT0083E: Could not get business detail for businessKey[<business key>].**

**Explanation:** The get business operation failed on the source registry.

**User Response:** Check the key exists in the source registry.

**UDUT0084E: Could not get binding detail for bindingKey[<binding key>].**

**Explanation:** The get binding operation failed on the source registry.

**User Response:** check the key exists in the source registry.

**UDUT0085E: Could not save tModel for tModelKey[<tModel key>].**

**Explanation:** The publish operation failed at the target registry.

**User Response:** Check the tModel is not referencing another entity (such as a tModel) that is not present in the target registry. This may occur if the 'importReferenced' property is set to false. Specify referenced tModels in the referencedtModels section of the Entity Definition File and set 'importReferenced' property in the configuration file to 'true'.

**UDUT0086E: Could not save business for businessKey[<business key>].**

**Explanation:** The publish operation failed at the target registry.

**User Response:** Check the businessEntity is not referencing another entity (such as a tModel) that is not present in the target registry. This may occur if the 'importReferenced' property is set to false. Specify referenced tModels in the referencedtModels section of the Entity Definition File and set 'importReferenced' property in the configuration file to 'true'.

**UDUT0087E: Could not save service for parent businessKey[<business key>].**

**Explanation:** The publish operation failed at the target registry.

**User Response:** Check the businessEntity specified as the parent of the businessService exists in the target registry.

**UDUT0088E: Could not save binding for parent serviceKey[<service key>].**

**Explanation:** The publish operation failed at the target registry.

**User Response:** Check the businessService specified as the parent of the bindingTemplate exists in the target directory.

**UDUT0089W: Did not save service for serviceKey[<service key>].**

**Explanation:** The parent business for the specified businessService does not exist.

**User Response:** Check the key value for the parent entity is correct in the Entity Definition File, and that the entity exists in the target registry.

**UDUT0090W: Did not save binding for bindingKey[<binding key>].**

**Explanation:** The parent service for the specified bindingTemplate does not exist.

**User Response:** Check the key value for the parent entity is correct in the Entity Definition File, and that the entity exists in the target registry.

**UDUT0091E: Could not delete business for businessKey[<business key>] because parent service did not exist.**

**Explanation:** The UDDI4J operation to delete the businessEntity with the specified key failed.

**User Response:** Check the userid and password property values in the configuration file and that the entity exists in the target UDDI registry.

**UDUT00921E: Could not delete service for serviceKey[<tModel key>].**

**Explanation:** The UDDI4J operation to delete the businessService with the specified key failed.

**User Response:** Check the userid and password property values in the configuration file and that the entity exists in the target UDDI registry.

**UDUT0093E: Could not delete binding for bindingKey[<binding key>].**

**Explanation:** The UDDI4J operation to delete the bindingTemplate with the specified key failed.

**User Response:** Check the userid and password property values in the configuration file and that the entity exists in the target UDDI registry.

**UDUT0094E: Could not delete tModel for tModelKey[<tModel key>].**

**Explanation:** The UDDI4J operation to delete the tModel with the specified key failed.

**User Response:** Check the userid and password property values in the configuration file and that the entity exists in the target UDDI registry.



**UDUT0096W: <entity type><key value> is not a valid UUID.**

**Explanation:** The key value entered does not comply with the format specified for a UUID in the UDDI specification.

**User Response:** Enter a valid UUID key.

**UDUT0097W: Did not save tModel for tModelKey[<tModel key>] as it already exists. Use the -overwrite argument to overwrite the tModel.**

**Explanation:** The tModel was not saved because the overwrite property is false

**User Response:** If the desired action is to overwrite existing entities, specify -overwrite on the command line or set the overwrite property in the configuration file to true.

**UDUT0098W: Did not save business for businessKey[<business key>] as it already exists. Use the -overwrite argument to overwrite the tModel.**

**Explanation:** The businessEntity was not saved because the overwrite property is false

**User Response:** If the desired action is to overwrite existing entities, specify -overwrite on the command line or set the overwrite property in the configuration file to true.

**UDUT0099W: Did not save service for serviceKey[<service key key>] as it already exists. Use the -overwrite argument to overwrite the tModel.**

**Explanation:** The businessService was not saved because the overwrite property is false

**User Response:** If the desired action is to overwrite existing entities, specify -overwrite on the command line or set the overwrite property in the configuration file to true.

**UDUT0100W: Did not save binding for bindingKey[<binding key key>] as it already exists. Use the -overwrite argument to overwrite the tModel.**

**Explanation:** The bindingTemplate was not saved because the overwrite property is false

**User Response:** If the desired action is to overwrite existing entities, specify -overwrite on the command line or set the overwrite property in the configuration file to true.

**UDUT0101W: Bad entity type: received<entity type>, expected<tModel|business|service|binding>.**

**Explanation:** The entered entity type was not recognized.

**User Response:** Specify arguments in accordance with the usage message.

**UDUT0102E: Promotion failed.**

**Explanation:** The promote function failed to complete.

**User Response:** Check the configuration properties file has correct settings.

**UDUT0106E: Unable to commit transaction.**

**Explanation:** The insertion of minimal entity data during the import function failed to commit to the database.

**User Response:** Check the database configuration. If necessary, turn on trace logging and look for the SQLException that is recorded.

**UDUT0107E: Unable to set auto-commit off on the database connection.**

**Explanation:** UDDI Utility Tools needs to control commits of data changes, however the attempt to turn off auto-commit failed.

**User Response:** Check the database configuration. If necessary, turn on trace logging and look for the SQLException that is recorded.

**UDUT0109E: The import function requires a UDDI entity definition file to be specified.**

**Explanation:** A required argument value was not specified.

**User Response:** Specify -definition <path to Entity Definition File> on the command line, or set the value of the UDDIEntityDefinitionFile property in the configuration file to the path to the Entity Definition File.

**UDUT0110E: A cyclic dependency exists in the referenced tModels. The reference from tModel with key [<tModel key>] to the tModel with key [<tModel key>] completes the detected cycle.**

**Explanation:** A cycle has been detected such that a tModel is being referenced by a tModel that it directly or indirectly references. This would cause the UDDI Utility Tools to enter an infinite loop trying to import referenced tModels, so the process is halted.

**User Response:** Edit the Entity Definition File and temporarily remove the reference to the tModel in the cycle, taking a note of the referenced details. After the import has successfully completed, update the tModel in the target registry to reintroduce the reference you previously removed. This can be done using the UDDI User Console, UDDI4J, or by creating a new Entity Definition File with just the tModel to be updated, and running the UDDI Utility Tools with the import function.

**UDUT0112E: An unexpected exception has occurred: <Exception message>.**

**Explanation:** An unexpected error occurred.

**User Response:** Check configuration file settings and all registries and databases are active. If necessary, contact your IBM support center.

**UDUT0113E: Could not get a response from UDDI registry at URL: <URL>.**

**Explanation:** A TransportException occurred while performing an UDDI4J operation on the UDDI registry at the specified URL.

**User Response:** Check configuration properties for the UDDI registry in question and ensure the UDDI registry is active.

**UDUT0114E: An IOException occurred trying to invoke 'java'.**

**Explanation:** When UDDI Utility Tools was invoked using the java -jar syntax, the invocation of the second JVM failed.

**User Response:** Check configuration property 'classpath' value is correct, and that Java is configured to run from the command line.

**UDUT0115I: Imported <entity count> entities and <referenced entity count> referenced entities.**

**Explanation:** Indicate that the import step of the import or promote function has completed, showing the number of entities imported.

**User Response:** None.

**UDUT0116W: Not all minimal entities could be removed. The following remain in the database:**

**Explanation:** A publish step was not successful which may have left one or more minimal entities in the target registry database. UDDI Utility Tools attempts to remove these minimal entities but in this case, the removal has failed. Following messages will indicate which minimal entities are left in the target registry.

**User Response:** You can attempt to remove the minimal entities using normal methods, such as the user console, UDDI4J, or using the delete function of the UDDI Utility Tools.

**UDUT0117W: Business minimal entities with businessKey [<business key>] has not been removed from the database.**

**Explanation:** A business minimal entity was orphaned in the target registry database and attempts to remove it failed.

**User Response:** Identify the orphaned minimal entity in the target and attempt to remove using normal UDDI methods, or by using the delete function of the UDDI Utility Tools.

**UDUT0118W: Service minimal entity with serviceKey [<service key>] has not been removed from the database.**

**Explanation:** A service minimal entity was orphaned in the target registry database and attempts to remove it failed.

**User Response:** Identify the orphaned minimal entity in the target registry and attempt to remove using normal UDDI delete methods, or by using the delete function of the UDDI Utility Tools.

**UDUT0119W: Binding Template minimal entity with bindingKey [<binding key>] has not been removed from the database.**

**Explanation:** A binding minimal entity was orphaned in the target registry database and attempts to remove it failed.

**User Response:** Identify the orphaned minimal entity in the target registry and attempt to remove using normal UDDI delete methods, or by using the delete function of the UDDI Utility Tools.

**UDUT0120W: TModel minimal entity with tModelKey [<tModel key>] has not been removed from the database.**

**Explanation:** A tModel minimal entity was orphaned in the target registry database and attempts to remove it failed.



**User Response:** Identify the orphaned minimal entity in the target registry and attempt to remove using normal UDDI delete methods, or by using the delete function of the UDDI Utility Tools.

**UDUT0121I: Created business minimal entity with businessKey [<business key>].**

**Explanation:** Indicates the minimal data required for a businessEntity has successfully been inserted in the target UDDI registry database.

**User Response:** None.

**UDUT0122I: Created service minimal entity with serviceKey [<service key>].**

**Explanation:** Indicates the minimal data required for a businessService has successfully been inserted in the target UDDI registry database.

**User Response:** None.

**UDUT0123I: Created binding template minimal entity with bindingKey [<binding key>].**

**Explanation:** Indicates the minimal data required for a bindingTemplate has successfully been inserted in the target UDDI registry database.

**User Response:** None.

**UDUT0124I: Created tModel minimal entity with tModelKey [<tModel key>].**

**Explanation:** Indicates the minimal data required for a tModel has successfully been inserted in the target UDDI registry database.

**User Response:** None.

**UDUT0125I: Deleted business minimal entity with businessKey [<business key>].**

**Explanation:** Indicates the minimal data inserted for a businessEntity was successfully removed from the target UDDI registry database. This would normally happen after a publish operation has failed.

**User Response:** None.

**UDUT0126I: Deleted service minimal entity with serviceKey [<service key>].**

**Explanation:** Indicates the minimal data required for a businessService was successfully removed from the target UDDI registry database. This would normally happen after a publish operation has failed.

**User Response:** None.

**UDUT0127I: Deleted binding template minimal entity with bindingKey [<binding key>].**

**Explanation:** Indicates the minimal data required for a bindingTemplate was successfully removed from the target UDDI registry database. This would normally happen after a publish operation has failed.

**User Response:** None.

**UDUT0128I: Deleted tModel minimal entity with tModelKey [<tModel key>].**

**Explanation:** Indicates the minimal data required for a tModel was successfully removed from the target UDDI registry database. This would normally happen after a publish operation has failed.

**User Response:** None.

**UDUT0129E: Find related businesses failed.**

**Explanation:** The UDDI4J find related businesses operation did not complete.

**User Response:** Check the configuration properties for the source registry, such as fromInquiryURL.

**UDUT0130E: Find businesses failed.**

**Explanation:** The UDDI4J find businesses operation did not complete.

**User Response:** Check the configuration properties for the source registry, such as fromInquiryURL.

**UDUT0131E: Find services failed.**

**Explanation:** The UDDI4J find services operation did not complete.

**User Response:** Check the configuration properties for the source registry, such as fromInquiryURL.

**UDUT0132E: Find tModels failed.**

**Explanation:** The UDDI4J find tModels operation did not complete.

**User Response:** Check the configuration properties for the source registry, such as fromInquiryURL.

**UDUT0133E: Find bindings failed.**

**Explanation:** The UDDI4J find bindings operation did not complete.

**User Response:** Check the configuration properties for the source registry, such as fromInquiryURL.

**UDUT0134I: Performing inquiry request ...**

**Explanation:** Indicated the find operation for selecting keys has started.

**User Response:** None.

**UDUT0135I: Extracted keys from inquiry results.**

**Explanation:** Indicates the find operation to select keys has completed successfully.

**User Response:** None

## UDUU (Web Services UDDI) messages

There are no messages issued by this component.

---

## Running the UDDI samples

The UDDI samples, and documentation on how to use them, are available through the Web Services UDDI samples link on the Samples Central page of the IBM WebSphere Developer Domain Web site.

---

## Installation Verification Program (IVP)

There are some samples available on the WebSphere Developer Domain (WSDD) web site (at <http://www7b.software.ibm.com/wsdd/library/samples/AppServer.html>) that are intended to provide an optional Installation Verification test, or IVP, for the UDDI Registry component.

This topic describes how to run these installation verification programs (IVPs) to verify that the IBM UDDI Registry has been installed correctly.

There are two IVP SOAP samples: SOAPSampleIVPa and SOAPSampleIVPb. They are intended to verify the successful installation of the product, and should be used in conjunction with the UDDI Users Console (GUI). SOAPSampleIVPa saves some data to the registry which you can then find using the GUI. Finally you can delete the data by running SOAPSampleIVPb.

The IVP samples are installed into the same target directory as the other SOAP samples and they use the same XML files as the basic Java SOAP samples.

SOAPSampleIVPa saves three businesses, six services (2 per business) and three tModels. The data structures are very basic and consist only of a name. The keys returned by the save\_\* UDDI API calls are then written to a file, SOAPSamp1eIVPa.out. SOAPSampleIVPb then reads in these keys from the file to delete the saved data from the UDDI registry.

**Note:** Each time you run SOAPSampleIVPa, it overwrites the output file SOAPSamp1eIVPa.out so, if you wish to use SOAPSampleIVPb to delete the data, you must run this before you next run SOAPSampleIVPa.

**Note:** As supplied, the IVP programs are written to work on a system without authentication. It is possible to configure the IVPs to work with authentication (see

<http://www7b.software.ibm.com/wsdd/library/samples/AppServer.html>), however, if possible it is recommended you run them on a non-authenticated system.

### Steps for this task

Perform the following steps on the same system as the UDDI Registry:

1. Ensure that DB2 and the WebSphere Administrative Server are started.
2. Start the WebSphere Administrator's Console and ensure the default server is started and the UDDI Registry Application is started.
3. For SOAP samples to work, you need to ensure that the Client Developer Kit for Java is either the one shipped with IBM WebSphere Application Server or a later IBM Developer Kit for Java.:
  - For Windows - ensure that <WebSphere-install-dir>\java\bin is present in the PATH statement before any other Developer Kits for Java
  - For Unix Platforms - ensure that <WebSphere-install-dir>/java/bin is present in the PATH statement before any other Developer Kits for Java

**Note:** You **must** use the IBM WebSphere supplied Developer Kit for Java or a later level of the IBM Developer Kit for Java.

For Windows, the default system path can be set via **Control Panel ...-> Settings ...-> System ...-> Advanced Properties ...-> Environment Variables**

Alternatively, this can be accomplished just for the shell where you plan to run the samples by modifying the path within the shell:

- For Windows - set path=<WebSphere-install-dir>\java\bin;%path%
  - For Unix Platforms - export PATH=<WebSphere-install-dir>/java/bin:\$PATH
4. Copy the samples and \*.xml files to a directory
  5. Compile both SOAPSsampleIVPa and SOAPSsampleIVPb by entering (from a command prompt):

```
'javac SOAPSsampleIVPa.java'
```

and

```
'javac SOAPSsampleIVPb'
```

6. Run SOAPSsampleIVPa by entering 'java SOAPSsampleIVPa'. This should publish a number of businesses and services and technical models into the registry.
7. Start your Web browser on the same system as the UDDI Registry.
8. To display the UDDI GUI home page, enter the following URL:
  - <http://localhost:9080/uddigui>
9. On the find page, complete the following steps:
  - a. Select the business radio button
  - b. In the data entry field, enter % (the wild card symbol)
  - c. Click **Find**

You should get a results page returned with three businesses (mybusiness1, mybusiness2, and mybusiness3). This demonstrates that the API and the UDDI user console are working correctly.
10. To see the services that are available for a business, click the "Show Services" option next to the business.
11. To delete all of the IVP data, run SOAPSsampleIVPb (from the command prompt as before - by entering 'java SOAPSsampleIVPb')
12. On the find page, complete the following steps:
  - a. Select the business radio button
  - b. In the data entry field, enter % (the wild card symbol)
  - c. Click **Find**

You should get an empty results page returned.

---

## Reporting problems with the IBM WebSphere UDDI Registry

If you report a problem with the IBM WebSphere UDDI Registry component to IBM, supply the following information:

1. A detailed description of the problem.
2. The build date and time of the version you are using. This can be obtained as follows:
  - In the installedApps subdirectory of the WebSphere installation location, you will find a subdirectory called UDDI\_Registry.<nodename>.<servername>.ear, where <nodename> is the name of the node into which the UDDI Registry application is installed, and <servername> is the name of the server. Within that subdirectory, you will find a file called version.txt. Include the contents of this file as part of your information.
  - If the UDDI Registry has been started with tracing enabled for the UDDI component, you should find a trace entry in the WebSphere trace log that includes the strings "getUDDIMessageLogger" and "UDDI Build : " followed by the build date and time, and the build system. Also include this information.
3. Any relevant log files and trace files.
  - If the problem occurred while setting up and installing the UDDI Registry application using one of the setup scripts, setupuddi.jacl or appserversetupuddi.jacl, supply the log output from running the script. (If you did not redirect the output from the script file to a log file, rerun the script, this time redirecting the output as described in the section 'Installing and Setting up a UDDI Registry'.) The log file is written to the directory from which you ran the setup script.
  - If the problem occurred while removing the UDDI Registry application using one of the remove scripts, removeuddi.jacl or appserverremoveuddi.jacl, supply the log output from running the script. (If you did not redirect the output from the script file to a log file, rerun the script, this time redirecting the output as described in the section 'Removing the UDDI Registry from a deployment manager cell' or 'Removing the UDDI Registry application from a single appserver'.) The log file is written to the directory from which you ran the remove script.
  - If the problem occurred while creating the UDDI Registry database using the UDDI DB2 Setup Wizard, supply the log file UDDIloadDB.log, which is written to the directory from which the wizard was run.
  - If the problem occurred while running the UDDI Registry, enable UDDI tracing (if not already enabled) and supply the trace log from the logs directory of the application server on which the UDDI Registry was running. See 'Turning on UDDI Trace' for details on how to enable UDDI tracing.
  - Also supply the WebSphere log files system.out and system.err.
  - Supply details of the version of IBM WebSphere Application Server you are running by executing the command versioninfo (Windows) or versioninfo.sh (Unix platforms) on both the application server and deployment manager nodes and directing the output to a log file.
4. If appropriate, any application code that you are using and the output produced by the application code.

In addition to the above, it is useful to run the WebSphere collector tool and send the resulting jar file(s) (two files if run from base application server AND DeploymentManager) to IBM. See "Running the collector tool" in the Information Center.

---

## Feedback

See the section on "Obtaining help from IBM" in the Information Center for details on seeking assistance.



---

## Chapter 12. Web services gateway: Enabling Web services

You use the Web services gateway to handle Web service invocations between Internet and intranet environments. You use it to make your internal Web services available externally, and to make external Web services available to your internal systems.

With the Web services gateway you administer the following resources:

- Web services.
- Channels.
- Filters.
- **5.1+** JAX-RPC handlers.
- Security bindings.
- UDDI references.

To make the Web services gateway files available for you to install, you select the gateway option when you install WebSphere Application Server (either by accepting the “typical install” option, or by choosing the “custom install” option **Web services > Web services gateway**). Note that when you install WebSphere Application Server, or an upgrade or fix pack, the installation program does *not* install the gateway. The installation program copies the gateway files into directories within your application server directory structure, then you install the gateway into your application server as a separate task.

When you install the gateway into your application server, any previously configured gateway is replaced with an upgraded but empty gateway. If you want to preserve an existing gateway configuration, you need to save the configuration before you install the upgrade, then restore the configuration after the upgrade is installed. To save your gateway configuration before you complete the installation of an upgraded gateway, see Preserving an existing gateway configuration.

To install the gateway into your application server, see Completing the installation.

Detailed instructions on how to enable Web services through the gateway are provided in the following tasks:

- Backing up and restoring a gateway configuration.
- Administering the Web services gateway.
- Running the Web services gateway samples.
- Passing SOAP messages with attachments through the Web services gateway.
- Developing Web services gateway extensions.
- Scaling the Web services gateway.
- Securing the Web services gateway.
- Tuning the Web services gateway.
- Troubleshooting the Web services gateway.

For a brief overview of the gateway and how it works, see [Web services gateway - frequently asked questions](#).

For a list of the major changes since the AlphaWorks preview version of the gateway, see [Web services gateway - What is new in this release](#).

For additional technical details of the Web services gateway, see the [gateway Javadoc](#).

For more information about working with Web services, visit the Web sites referenced in [Web services: Resources for Learning](#).

---

## Web services gateway - frequently asked questions

This topic provides answers to a set of frequently asked questions about the Web services gateway.

- What are Web services?
- What is the IBM Web services gateway?
- How does the Web services gateway work?
- What problems are solved by the Web services gateway?
- Who should use the Web services gateway?
- What is the difference between the Apache SOAP channel and the SOAP over HTTP channel?

### What are Web services?

Web services are modular applications that interact with one another across the Internet. Web services are based on shared, open and emerging technology standards and protocols (such as SOAP, UDDI, and WSDL) and can communicate, interact, and integrate with other applications, no matter how they are implemented.

### What is the IBM Web services gateway?

The gateway is a middleware component that bridges the gap between Internet and intranet environments during Web service invocations. You use the gateway to manage:

- Web services.
- Channels that carry requests to and responses from the services.
- Filters that act upon the services.
- **5.1 +** JAX-RPC handlers that act upon the services.
- References to UDDI registries in which services can be registered.
- Security bindings.

### How does the Web services gateway work?

The gateway builds upon the Web Services Description Language (WSDL) and the Web Services Invocation Framework (WSIF) for deployment and invocation.

You deploy a Web service to the gateway by deploying a WSDL file that describes how the gateway accesses the service. You deploy the WSDL file to a UDDI registry or to a Web address. You can send requests through the gateway to a Java class, an enterprise bean, a SOAP server or a SOAP over JMS server (including another gateway).

A request to the Web services gateway arrives through a channel, is translated into an internal form, then passes through any filters that are registered for the requested service, and finally sent on to the service implementation. Responses follow the same path in reverse.

### What problems are solved by the Web services gateway?

The gateway solves the following problems:

- **Securely “externalizing” Web services:** Business applications that are exposed as Web services can be used by any Web service-enabled tool, regardless of the implementation details. To better integrate your business processes, you might want to expose these assets to business partners, customers and suppliers who are outside the firewall. The gateway lets clients from outside the firewall use Web services that are hosted within your enterprise. Using the gateway, you can control access to each of these services.
- **Better return on investment:** Any number of partners can reuse a process that you develop as a Web service.
- **Use of existing infrastructure:** With the gateway, you can use your existing messaging infrastructure to make Web service requests, and use your existing Web services for external process integration.



- **Protocol transformation:** You might use one particular messaging protocol to invoke Web services, while your partners use some other protocol. Using the Web services gateway, you can trap the request from the client and transform it to another messaging protocol.

### Who should use the Web services gateway?

Any enterprise that chooses to share its resources selectively with its business partners and customers. IT managers and developers, who deploy resources, can also benefit from this technology.

### What is the difference between the Apache SOAP channel and the SOAP over HTTP channel?

The Apache SOAP channel and the SOAP over HTTP channel both support SOAP applications that are SOAP 1.1 compatible (for example Apache SOAP 2.3 and Axis SOAP 1.0). Therefore if you have an application that uses a production-supported Axis 1.0 SOAP stack that is generating SOAP 1.1, then this application can use either channel. However the Apache SOAP channel is deprecated and wholly superseded by the SOAP over HTTP channel.

If you use the Apache SOAP channel, then the SOAP message format must be RPC style. The SOAP over HTTP channel has all the functionality of the Apache SOAP channel, and also supports:

- Document style SOAP message format.
- SOAP messages with attachments.
- **5.1+** Java API for XML-based Remote Procedure Call (JAX-RPC) handlers.
- **5.1+** Selective SOAP parsing.

---

## Web services gateway - What is new in this release

The Web services gateway was first released on AlphaWorks on 21 December 2001. Here is a list of the main differences between the AlphaWorks edition and this version:

- The gateway is now built using enterprise beans. This new construction supports the Web services gateway only in an application server that has an EJB container. The gateway no longer runs in the Tomcat server.
- The gateway includes UDDI integration, therefore you can deploy and remove Web services to a UDDI registry as well as to a Web address.
- In addition to the Apache SOAP channel, the following channels now exist:
  - A SOAP over HTTP channel.
 

This channel supersedes the Apache SOAP channel. It has all the functionality of the Apache SOAP channel, and also supports:

    - RPC style and Document style SOAP message formats.
    - SOAP messages with attachments.
    - **5.1+** Java API for XML-based Remote Procedure Call (JAX-RPC) handlers.
    - **5.1+** Selective SOAP parsing.
  - **5.1+** A synchronous SOAP over Java Messaging Service (JMS) channel.
- The gateway supports bidirectional interactions (that is, both inbound and outbound requests) directly, by deploying two instances of each type of channel. To achieve this configuration with the AlphaWorks version, you deployed two instances of the Web services gateway; one for inbound communication and one for outbound communication.

### Note:

- Interceptors have been renamed as filters.
- Channels, filters and UDDI references are deployed to the Web services gateway, then associated with individual Web services. When you configure a Web service, you choose the following entities:
  - The channels on which the Web service is available.
  - Any filters that apply.

- **5.1+** Any JAX-RPC handlers that apply.
- Any UDDI references to which the Web service is deployed.
- Any security bindings that apply.
- You can change the channels, filters, handlers and UDDI references that are associated with a deployed service without removing the service.
- You can deploy multiple targets for a single service (that is, more than one implementation of a service that has the same service interface).
- You can back up and restore a gateway configuration.
- You can scale the Web services gateway.
- You can configure the gateway for secure transmission of SOAP messages using tokens, keys, signatures and encryption in accordance with the emerging Web Services Security (WS-Security) specification.
- You can set security (basic authorization) on the individual methods of a Web service, and on the whole gateway.
- The gateway can invoke Web services over HTTPS.
- You can configure the gateway to work with authenticating proxies.

---

## Web services gateway - Completing the installation

This task assumes that when you installed WebSphere Application Server you selected the Web services gateway either by accepting the “typical install” option, or by choosing the “custom install” option **Web services > Web services gateway**.

When you install WebSphere Application Server, or apply an upgrade or fix pack, the installation program does not install the gateway enterprise applications. When you run the installation program, and select the Web services gateway component, the files that are required for running the gateway are copied into directories under *WebSphere\_DeployMgr\_root*, where *WebSphere\_DeployMgr\_root* is the deployment manager root directory (by default WebSphere/DeploymentManager ). For example, if you install WebSphere Application Server and accept the default directory names, then the location of the `installableApps` directory is `WebSphere/DeploymentManager/installableApps` .

When you install the gateway into your application server (by completing the steps described in this topic), any previously configured gateway is replaced with an upgraded but empty gateway. If you want to preserve an existing gateway configuration, you need to save the configuration before you install the upgraded gateway, then restore the configuration after the upgrade is installed. To save your gateway configuration before you complete the installation of an upgraded gateway, see [Preserving an existing gateway configuration](#).

The following table lists the gateway files, and the locations into which they are placed. The **Location** column shows the subdirectory under *WebSphere\_DeployMgr\_root*.

File name	Purpose	Location
wsgw.ear	The Web services gateway application	/installableApps
wsgwsoap1.ear	The Apache SOAP channel 1 application	/installableApps
wsgwsoap2.ear	The Apache SOAP channel 2 application	/installableApps
wsgwsoaphttp1.ear	The SOAP over HTTP channel 1 application	/installableApps
wsgwsoaphttp2.ear	The SOAP over HTTP channel 2 application	/installableApps

File name	Purpose	Location
wsgwsoapjms1.ear	The synchronous SOAP over Java Messaging Service (JMS) channel 1 application	/installableApps
wsgwsoapjms2.ear	The synchronous SOAP over JMS channel 2 application	/installableApps
wsgwauth.ear	The Web service operation-level security application	/installableApps
WSGWResourceBundles.jar	System messages for the Web services gateway	/lib
Installation scripts	Installation of the Web services gateway	/WSGW/scripts/install
Authorization scripts	Generation of authorization beans for Web service operation-level security	/WSGW/scripts/auth

To complete the gateway installation, you have two choices:

- Install into an application server that is part of a deployment manager cell.
- Install into a stand-alone application server (that is, an application server that is not part of a deployment manager cell).

The gateway functions identically in either environment, but if you install the gateway into an application server that is part of a deployment manager cell, then you can use the deployment manager tooling for that cell to complete the gateway installation on any server within the cell, irrespective of where the server is physically hosted. Note however that the gateway is not fully integrated into the WebSphere Application Server Network Deployment tooling:

- You cannot use the full power of the deployment manager tooling to automatically install or update a set of gateways across all application servers in a cell. You must follow the more manual steps for each server that are described in this topic.
- You cannot install the gateway into an application server that is part of a WebSphere Application Server Network Deployment cluster. See "Setting up a multinode environment" in the Information Center.

To achieve high availability and load balancing, follow the steps described in *Scaling the Web services gateway*. The gateways that you include in your load-balancing group can be hosted on stand-alone application servers, or application servers that are members of a deployment manager cell, or a mixture of the two.

To finish the Web services gateway installation, complete the following steps:

1. Confirm that your system configuration complies with the Web services gateway prerequisites and constraints.
2. Optional: Preserve an existing gateway configuration.
3. Optional: Enable security.
4. Either Install the gateway into a deployment manager cell, or Install the gateway into a stand-alone application server.
5. Test the installation.

## Web services gateway - prerequisites and constraints

WebSphere Application Server Network Deployment is not a stand-alone product for running enterprise applications. If you want to deploy the gateway you also need to install the base WebSphere Application Server product. Although it is not installed by default, a copy of the base WebSphere Application Server product is packaged with WebSphere Application Server Network Deployment.

The gateway is not fully integrated into the WebSphere Application Server Network Deployment tooling:

- When you install WebSphere Application Server, or an upgrade or fix pack, the installation program does *not* install the gateway. The installation program copies the gateway files into directories within your application server directory structure, then you install the gateway into your application server as a separate task.
- You cannot use the full power of the deployment manager tooling to automatically install or update a set of gateways across all application servers in a cell.
- You cannot install the gateway into an application server that is part of a WebSphere Application Server Network Deployment cluster. See "Setting up a multinode environment" in the Information Center

For more information about these restrictions, see [Completing the installation](#).

If you want to enable gateway-level security, you must do so before you install the gateway.

This version of the Web services gateway is also subject to the following constraints:

- WSDL definitions for target services must use XML Schema Version 2001. For more information, see [Troubleshooting the Web services gateway](#).
- Install the gateway application (wsgw.ear) before you install channel and filter applications. If you need to reinstall the gateway application, uninstall all channels and filters first, then reinstall them after the gateway application.
- If the WSDL for your service contains <soap:header> elements within the <wsdl:definition> element, then the gateway passes the SOAP headers through. This behavior is correct. However, you also see the following effects:
  - The SOAP headers are not included in the WSDL that the gateway generates.
  - If you set the "must understand" flag on the SOAP message, then you get an error message.
- The Apache SOAP channel is deprecated and wholly superseded by the SOAP over HTTP channel.
- If you are using the Apache SOAP channel, then the SOAP message format must be RPC style. To handle Document style SOAP messages, use the SOAP over HTTP channel, which supports both RPC style and Document style SOAP messages.

You might also find it useful to enable trace for all gateway components:

1. Set the WebSphere Application Server trace string to `com.ibm.wsgw.*=all=enabled`
2. Write trace, stdout and stderr for the application server to a well-known location.

For more information, see ["Enabling trace"](#) in the Information Center.

## Preserving an existing gateway configuration

When you install WebSphere Application Server, or an upgrade or fix pack, the installation program does *not* install the gateway. The installation program copies the gateway files into directories within your application server directory structure, then you install the gateway into your application server as a separate task.

When you install the gateway, any previously configured gateway is replaced with an upgraded but empty gateway. If you want to preserve an existing gateway configuration, you need to save the configuration before you install the upgraded gateway into your application server, then restore the configuration after the upgrade is installed.

To preserve your existing gateway configuration, complete the following steps:

1. Back up the existing configuration.
2. Complete the installation of the upgraded gateway.
3. Restore the previously-saved configuration.

## Installing the gateway into an application server that is part of a deployment manager cell

Use this task to install the gateway into an application server that is part of a node of an existing deployment manager cell.

This topic assumes that you have completed the pre-installation steps described in [Completing the installation](#) and that the gateway application files are therefore available in directories under *WebSphere\_DeployMgr\_root*, where *WebSphere\_DeployMgr\_root* is the deployment manager root directory (by default *WebSphere/DeploymentManager*). This topic also assumes that you have already created or selected an application server that is part of a node in a deployment manager cell, and that this application server is not also part of a cluster.

When you install the gateway into your application server (by completing the steps described in this topic), any previously configured gateway is replaced with an upgraded but empty gateway. If you want to preserve an existing gateway configuration, you need to save the configuration before you install the upgraded gateway, then restore the configuration after the upgrade is installed. To save your gateway configuration before you complete the installation of an upgraded gateway, see [Preserving an existing gateway configuration](#).

**5.1** When you run the `jacl` script to install the gateway upgrade (as described in this topic), the previous versions of the gateway applications are not automatically uninstalled. The `jacl` script attempts to uninstall the previous gateway applications before it installs the new ones, but the script assumes that the previous applications have the same installed application names as the new ones. However, the default names used for the installed gateway applications were changed between WebSphere Application Server Version 5.0.2 and Version 5.1. Therefore you must manually uninstall the previous versions of the gateway applications after you have saved the configuration (optional), and before you use the `jacl` script to install the upgrade.

To install the gateway into an application server that is part of a node in a deployment manager cell, complete the following steps:

1. Start the application server.
2. Clear your class path. You can use the following command:
  - (Windows systems): `set CLASSPATH=`
  - (UNIX and Linux systems): `unset CLASSPATH`
3. Navigate to the *WebSphere\_DeployMgr\_root*/WSGW/scripts/install directory.
4. Enter the command: *WebSphere\_DeployMgr\_root*/bin/wsadmin.ext -f setupWSGW.jacl *parm1* ... *parmN* where:
  - *parm1* is the *WebSphere\_DeployMgr\_root* directory
  - *parm2* is the server name
  - *parm3* is the node name (this is case sensitive)

For example (Windows systems, shown on multiple lines for publication):

```
C:\Progra~1\WebSphere\DeploymentManager\bin\wsadmin.bat -f setupWSGW.jacl
C:/Progra~1/WebSphere/DeploymentManager server1 PHJ2
```

**Note:** The use of forward slashes (/) is compulsory for this command, even on Windows systems.

**Tip:** If you create a symbolic link to the `wsadmin.sh` file, then you need not type the directory path to the `wsadmin.sh` file in the previous command. To create this symbolic link, issue the following command from within the *WebSphere\_DeployMgr\_root*/WSGW/scripts/install directory:

```
ln -s WebSphere/V5R1M0/DeploymentManager/bin/wsadmin.sh wsadmin.sh
```

After you run the `setupWSGW.jacl` command, system messages are displayed confirming that the following initial set of gateway applications have installed successfully:

- The gateway (wsgw.ear).
- Apache SOAP channel 1 (wsgwsoap1.ear).
- SOAP over HTTP channel 1 (wsgwsoaphttp1.ear).

**Note:** **5.1+** In some circumstances, the “installed successfully” system messages are immediately followed by error messages saying that these same applications are *not* installed. These error messages should not appear, and you should ignore them, because the installation of the applications has succeeded.

5. Install additional gateway applications. For example:
  - Apache SOAP channel 2 (wsgwsoap2.ear)
  - SOAP over HTTP channel 2 (wsgwsoaphttp2.ear)

To install additional gateway applications, complete the following steps:

**Note:**

- If you prefer, you can install these EAR files using the WebSphere Application Server administrative console, as described in the final step of Enabling operation-level authorization.
- **5.1+** Before you install the synchronous SOAP over JMS channel, see the additional information in Installing the synchronous SOAP over JMS channel.

- a. Navigate to the *WebSphere\_DeployMgr\_root/bin* directory.
- b. Start the application server.
- c. Start the WebSphere administration program (by entering the command *wsadmin.ext*).
- d. For each additional Web services gateway enterprise application that you install, enter the following commands at the *wsadmin>* prompt:

```
$AdminApp install path_to_ear_file {-appname application
 -server your_server -node your_node_name}
$AdminConfig save
```

where:

- *application* is the name of the enterprise application
- *path\_to\_ear\_file* is the name and location of the enterprise application EAR file
- *your\_node\_name* is the node name (this value is case sensitive)

For example (Windows systems):

```
$AdminApp install
 C:/Progra~1/WebSphere/DeploymentManager/installableApps/wsgwsoap2.ear
 {-appname wsgwsoap2 -server server1 -node PHJ2}
$AdminConfig save
```

- e. After you install all your additional Web services gateway enterprise applications, close the WebSphere administration program by entering *quit* or *exit* at the *wsadmin>* prompt.

6. Stop then restart the application server.

If you want more than one gateway installation (for example, to scale the Web services gateway), repeat the steps in this topic for another application server that is part of a node in a deployment manager cell.

You are now ready to test the installation. Run the test on every application server on which you installed the gateway.

## Installing the gateway into a stand-alone application server

Use this task to install the gateway into an individual application server instance in your network space.

A stand-alone application server is an application server that is not part of a deployment manager cell. Before you install the gateway into a stand-alone application server, you should consider whether you might subsequently want to add the machine on which the application server is hosted as a node of a cell.



If you do subsequently add the machine to a cell, then the installed gateway application will be deleted by the node synchronization process. You cannot install the gateway into a stand-alone application server that is part of a cluster. See "Setting up a multinode environment" in the Information Center.

When you install the gateway into your application server (by completing the steps described in this topic), any previously configured gateway is replaced with an upgraded but empty gateway. If you want to preserve an existing gateway configuration, you need to save the configuration before you install the upgraded gateway, then restore the configuration after the upgrade is installed. To save your gateway configuration before you complete the installation of an upgraded gateway, see Preserving an existing gateway configuration.

**5.1** When you run the `jacl` script to install the gateway upgrade (as described in this topic), the previous versions of the gateway applications are not automatically uninstalled. The `jacl` script attempts to uninstall the previous gateway applications before it installs the new ones, but the script assumes that the previous applications have the same installed application names as the new ones. However, the default names used for the installed gateway applications were changed between WebSphere Application Server Version 5.0.2 and Version 5.1. Therefore you must manually uninstall the previous versions of the gateway applications after you have saved the configuration (optional), and before you use the `jacl` script to install the upgrade.

This topic assumes that you have completed the pre-installation steps described in Completing the installation and that the gateway application files are therefore available in directories under `WebSphere_DeployMgr_root`, where `WebSphere_DeployMgr_root` is the deployment manager root directory (by default `WebSphere/DeploymentManager`). Before you can install and run the gateway in a stand-alone application server, you must copy these files over to the application server by completing the following steps:

- Stop the application server into which you plan to install the Web services gateway.
- Copy all the EAR files with file names that begin `wsgw` from the `WebSphere_DeployMgr_root/installableApps` directory of the machine on which you installed WebSphere Application Server Network Deployment into the `stand-alone_install_root/installableApps` directory of the target application server install tree, where `stand-alone_install_root` is the root directory for the installation of your target application server (by default `WebSphere/AppServer`).
- Copy the `WebSphere_DeployMgr_root/lib/WSGWResourceBundles.jar` file into the `stand-alone_install_root/lib` directory.
- Copy the `WebSphere_DeployMgr_root/WSGW` directory and all files and directories within it into the `stand-alone_install_root/WSGW` directory.

To install the gateway into your stand-alone application server, complete the following steps:

1. Start the application server.
2. Clear your class path. You can use the following command:
  - (Windows systems): `set CLASSPATH=`
  - (UNIX and Linux systems): `unset CLASSPATH`
3. Navigate to the `stand-alone_install_root/WSGW/scripts/install` directory.
4. Enter the command: `stand-alone_install_root/bin/wsadmin.ext -f setupWSGW.jacl parm1 ... parmN` where
  - `parm1` is the `stand-alone_install_root` directory
  - `parm2` is the server name
  - `parm3` is the node name (this is case sensitive)

For example (Windows systems, shown on multiple lines for publication):

```
C:\Progra~1\WebSphere\AppServer\bin\wsadmin.bat -f setupWSGW.jacl
C:/Progra~1/WebSphere/AppServer server1 PHJ2
```



**Note:** The use of forward slashes (/) is compulsory for this command, even on Windows systems. When you run the `setupWSGW.jac1` command, the following initial set of gateway applications is installed:

- The gateway (`wsgw.ear`).
  - Apache SOAP channel 1 (`wsgwsoap1.ear`).
  - SOAP over HTTP channel 1 (`wsgwsoaphttp1.ear`).
5. Install additional gateway applications. For example:
- Apache SOAP channel 2 (`wsgwsoap2.ear`)
  - SOAP over HTTP channel 2 (`wsgwsoaphttp2.ear`)

To install additional gateway applications, complete the following steps:

**Note:**

- If you prefer, you can install these EAR files using the WebSphere Application Server administrative console, as described in the final step of Enabling operation-level authorization.
  - **5.1+** Before you install the synchronous SOAP over JMS channel, see the additional information in Installing the synchronous SOAP over JMS channel.
- a. Navigate to the `stand-alone_install_root/bin` directory.
  - b. Start the application server.
  - c. Start the WebSphere administration program (by entering the command `wsadmin.ext`).
  - d. For each additional Web services gateway enterprise application that you install, enter the following commands at the `wsadmin>` prompt:

```
$AdminApp install path_to_ear_file {-appname application
 -server your_server -node your_node_name}
$AdminConfig save
```

where

- *application* is the name of the enterprise application
- *path\_to\_ear\_file* is the name and location of the enterprise application EAR file
- *your\_node\_name* is the node name (this value is case sensitive)

For example (Windows systems):

```
$AdminApp install
 C:/Progra~1/WebSphere/AppServer/installableApps/wsgwsoap2.ear
 {-appname wsgwsoap2 -server server1 -node PHJ2}
$AdminConfig save
```

- e. After you install all your additional Web services gateway enterprise applications, close the WebSphere administration program by entering `quit` or `exit` at the `wsadmin>` prompt.
6. Stop then restart the application server.

If you want more than one gateway installation (for example, to scale the Web services gateway), repeat the steps provided in this topic for another stand-alone application server.

You are now ready to test the installation. Run the test on every application server on which you installed the gateway.

## Testing the Web services gateway installation

Use this task to test that the Web services gateway is installed correctly.

This topic assumes that you have completed the steps described in Completing the installation.

To test the basic installation of the Web services gateway, complete the following steps:

1. In a Web browser, go to `http://host:port/wsgw` where *host* and *port* are the host name and port number on which your HTTP server is listening. The browser displays the following message:

## IBM Web services gateway

What do you want to do?

- [Run the admin client](#)
  - [View the product ID](#)
2. Test the Apache SOAP channels by using your Web browser to display the Web page at `http://host:port/wsgwengine/soaprpcrouter`, where *engine* is either *soap1* or *soap2*. The browser displays the following message: Sorry, I don't speak via HTTP GET - you have to use HTTP POST to talk to me.
  3. Test the SOAP over HTTP channels by using your Web browser to display the Web page at `http://host:port/wsgwengine/soaphttpengine`, where *engine* is either *soaphttp1* or *soaphttp2*. The browser displays the message: And now... Some Services, followed by a list of the Web services that are currently deployed to the channel.

If you do not see these messages, your server is not configured correctly - in which case, see [Troubleshooting the Web services gateway](#).

---

## Backing up and restoring a gateway configuration

There are two situations in which you might want to back up a gateway configuration:

- To preserve an existing gateway configuration.
- To scale the Web services gateway.

The **Back up** options save the deployment details for all the Web services and associated resources (channels, filters, UDDI references and so on) that are currently deployed to the gateway. There are two backup options: **Save Gateway Configuration > Private** and **Save Gateway Configuration > Shared**. For an explanation of these terms, and detailed steps for this task, see [Backing up a gateway configuration](#).

The **Restore** option uses a previously-saved gateway configuration to populate an empty installation of the gateway with the same deployment details. For detailed steps for this task, see [Restoring a gateway configuration](#).

## Backing up and restoring UDDI publication links

When you save a gateway configuration, the deployment details through which a deployed Web service has been published to UDDI by the gateway are not saved. Here are the reasons for not backing up these details:

- When a gateway is upgraded, or otherwise removed without first using the gateway administrative console to un-publish each Web service from UDDI, then the associated entries in UDDI registries are not removed.
- When a gateway is restored, if the UDDI publication details are also restored, then the Web services are re-published to UDDI by the gateway.
- When a gateway publishes the same service twice to a UDDI registry, the registry does not overwrite the initial publication. The registry creates a second copy of the TModel and Service Definition.
- When a gateway removes a service from a UDDI registry, it only removes the last one published. This means that it can leave behind outdated TModels and Service Definitions.

To preserve the integrity of the UDDI registries, restoring the gateway does not automatically re-publish any Web services to UDDI. But if you know that a service has been un-published from UDDI (for example because you removed it manually before backing up the gateway configuration) then you can safely re-publish it to UDDI after the configuration is restored.

## Backing up a gateway configuration

The **Back up** options write out, to a single file, the deployment details for all the Web services and associated resources (channels, filters, UDDI references and so on) that are currently deployed to the gateway.

There are two situations in which you might want to back up a gateway configuration:

- To preserve an existing gateway configuration.
- To scale the Web services gateway.

When you install a gateway upgrade, any previously configured gateway is replaced with an upgraded but empty gateway. Use the **Save Gateway Configuration > Private** option to preserve an existing gateway configuration before you complete the installation of an upgraded gateway. This saved configuration contains deployment details (for example URIs) that are specific to a given instance of the gateway, so you should only use it to save a configuration that you intend to re-apply to the same gateway instance.

When you want to scale the gateway, use the **Save Gateway Configuration > Shared** option to save a shareable version of your gateway configuration. This saved configuration does not include machine-specific values, and therefore can be shared by all the gateways in your scaled gateway deployment group.

To back up your gateway configuration, complete the following steps:

1. Display the Web services gateway administrative user interface.
2. Optional: If you are backing up to preserve your gateway configuration before you install a gateway upgrade, and any of your deployed Web services are also published by the gateway to UDDI registries, use the Listing and managing gateway-deployed Web services option to un-publish them from UDDI. Make a note of the UDDI deployment details so that you can re-publish them to UDDI after you restore the gateway configuration (the reasons why this manual step is necessary are explained in Backing up and restoring UDDI publication links).
3. In the gateway administrative user interface navigation pane, click the following link:  
**Gateway**
  - [Back Up](#)The Save Gateway Configuration window opens.
4. In the **Location** field, type the path to your configuration file. The path to your configuration file must point to a local drive on the machine on which the gateway is currently running. The file name for the configuration file can be any valid Java file name.
5. If you are backing up to preserve the configuration of a gateway before you install a gateway upgrade, select **Private**. If you are backing up to create a scaled gateway deployment, select **Shared**.
6. Click **OK**.

Your gateway configuration is saved to the location you specified.

## Restoring a gateway configuration

This topic assumes that you have already completed the following tasks:

1. Back up your gateway configuration.
2. Complete the installation of an upgraded gateway.

The **Restore** option uses the information contained in a previously-saved file to populate an empty installation of the gateway with the same deployment details for those Web services and associated resources.

The **Restore** option automatically detects whether the configuration file that it is restoring contains a **Private** or a **Shared** gateway configuration (for an explanation of these terms, see Backing up a gateway configuration).

To restore a gateway configuration, complete the following steps:

1. Check that the target gateway is empty. Remove any Web services and associated resources that are deployed to the gateway. If you have just completed the installation of an upgraded gateway, then the installed gateway is empty.
2. In the gateway administrative user interface navigation pane, click the following link:

**Gateway**

- [Restore](#)

The Restore Gateway Configuration window opens.

3. In the **Location** field, type the path to your configuration file.

The path to your configuration file must point to a local drive on the machine on which the gateway is currently running. The saved file is not specific to a given version of the gateway.

If you are restoring a file that was backed up using the **Save Gateway Configuration > Private** option, then the backup contains deployment details (for example URIs) that are specific to a given gateway instance. If you restore this type of backup file to a different instance (for example a gateway on a host with a different network identity), then the restore succeeds but you have to use the gateway administrative user interface to correct any errors in the deployment details.

4. Click **OK**.

Your gateway configuration is restored to the gateway that you are currently administering.

If you have just restored a configuration after installing an upgraded gateway, and any of the deployed Web services were previously published to UDDI registries, use the Listing and managing gateway-deployed Web services option (and the notes you made when you backed up the configuration) to re-publish these Web services to UDDI. The reasons why this manual step is necessary are explained in Backing up and restoring UDDI publication links.

---

## Administering the Web services gateway

If you have just completed the installation of an upgraded gateway, your gateway configuration is currently empty. If you want to restore a saved configuration, you must do so before you add any new channels, filters, handlers, UDDI references or Web services.

You should not make changes to a gateway that is part of an unlocked configuration load-balancing group unless you are making changes, under controlled conditions, to all members of the group.

To administer the Web services gateway, complete the following steps:

1. Start the WebSphere Application Server Administrative Server.
2. Open the following Web page: `http://host:port/wsgw/admin/index.html` where *host* and *port* are the host name and port number on which your HTTP server is listening. For example, `localhost:8080` or `localhost:9080`. The main administration page is displayed:

# IBM Web Services Gateway

## Gateway

- **Configure**
- **Security**
- **Back Up**
- **Restore**
- **About**

## Services

- **List**
- **Deploy**
- **Remove**

## Channels

- **List**
- **Deploy**
- **Remove**

## Filters

- **List**
- **Deploy**
- **Remove**

## Handlers

- **List**
- **Deploy**
- **Remove**

## UDDI References

- **List**
- **Deploy**
- **Remove**

## Welcome to the IBM Web Services Gateway systems administration console

To begin, please make a selection from the options to the left

To modify features after initial deployment, first select "list" under the feature and then select the item you want to modify

The gateway administration options are listed in the left-hand panel. The order in which these options are listed is significant, for the following reasons:

- If you change the namespace URI or WSDL URI (using the **Gateway > Configure** option), you break the link back to the gateway for every Web service that you have already deployed. You must therefore set these URIs before you deploy any Web services.
- When you deploy a Web service to the gateway, you choose the following entities:
  - The channels on which it is available.
  - Any filters that apply to it.
  - **5.1+** Any JAX-RPC handlers that apply to it.
  - Any UDDI references to which it is deployed.

Each of these choices is made from a list of resources that are already deployed to the gateway. Therefore you might want to deploy your channels, filters, handlers and UDDI references to the gateway before you deploy the Web services that use them.

3. Use the gateway administration options to configure the gateway as described in the following topics:
  - Setting the namespace URI and WSDL URI for the Web services gateway
  - Working with channels
  - Working with filters
  - **5.1+** Working with JAX-RPC handlers
  - Working with UDDI references
  - Working with Web services

In all of the gateway administration panels, fields marked with asterisks are required. After you deploy a channel, filter, handler or UDDI reference, refresh all other open browser windows to ensure that current lists are displayed.

## Setting the namespace URI and WSDL URI for the Web services gateway

Initial values for the namespace URI and WSDL URI are automatically configured when you install the Web services gateway.

When you change these URIs, you break the link back to the Web services gateway for every Web service that you have already deployed, so you must set these URIs before you deploy any Web services to the gateway.

You should not make changes to these URIs for a gateway that is part of an unlocked configuration load-balancing group.

When you deploy a Web service to the gateway, these two URIs are used in the following ways:

- The **Namespace URI for services** is used as the namespace for the gateway services in exported WSDL documents.
- The **WSDL URI for exported definitions** is used to generate the Web address in import statements within exported WSDL documents.

To set the namespace URI and WSDL URI for the Web services gateway, complete the following steps:

1. Display the Web services gateway administrative user interface.
2. In the navigation pane, click the following link:

### **Gateway**

- [Configure](#)

The gateway configuration form is displayed:

# IBM Web Services Gateway

## Gateway

- **Configure**
- Security
- Back Up
- Restore
- About

## Services

- List
- Deploy
- Remove

## Channels

- List
- Deploy
- Remove

## Filters

- List
- Deploy
- Remove

## Handlers

- List
- Deploy
- Remove

## UDDI References

- List
- Deploy
- Remove

## Configure Gateway

Namespace URI for services

WSDL URI for exported definitions

## Proxy Configuration

Enable proxy authentication

Proxy user

Proxy password

Use Gateway proxy credentials for invoking WebServices

3. In the **Namespace URI for services** field, type the new name.

There is no fixed syntax for the namespace URI, but whatever name you choose is likely to be more effective if it observes the following guidelines:

- The name begins with urn:

See the guidance on Internet standards for the syntax of Uniform Resource Names (URNs) at <http://www.ietf.org/rfc/rfc2141.txt>.

- The name is globally unique.
- The name reflects your company name.

4. In the **WSDL URI for exported definitions** field, type the new name.



The initial value is the gateway best guess at the right value, but you probably want to overwrite it with a new value. For example it might guess a local URI such as `http://h1dswrth:9080/wsgw`, and because you are giving the WSDL to people in other companies, you modify this value to `http://h1dswrth.your_company.com/wsgw`. Note that only the *host* and *port* parts of the initial value are modified. This URI must always start with `http://` and end with `/wsgw`.

5. Click **Apply Changes**.

You are now ready to deploy your channels, filters, handlers and UDDI references to the gateway.

You might also want to enable proxy authentication for the gateway.

## Working with channels

Before you can work with a channel, you must install the channel application in WebSphere Application Server as described in *Installing the gateway-supplied channels*.

You should not make changes to the channel deployment settings for a gateway that is part of an unlocked configuration load-balancing group unless you are making changes, under controlled conditions, to all members of the group.

Two versions of each type of channel are supplied so that, for each channel type, you can set up separate channels for inbound and outbound requests. For more information, see *Channels - entry points to the Web services gateway*.

From the navigation pane of the Web services gateway administrative user interface, you can choose the following actions for **Channels**:

- **List** to list the deployed channels, and modify their deployment details.
- **Deploy** to deploy a channel.
- **Remove** to remove channels.

### Channels - Entry points to the Web services gateway

Channels form entry points to the Web services gateway and carry requests and responses between Web services and the Web services gateway.

A request to the Web services gateway arrives through a channel, is translated into a WSIF message, then passed through any filters that are registered for the requested service, and finally sent on to the service implementation. Responses follow the same path in reverse.

Before you can use a channel, you must install the channel application in WebSphere Application Server then deploy the channel to the Web services gateway. A deployed channel is not used until you deploy a Web service that uses the channel.

Two versions of each type of channel are supplied with the gateway. This is so that, for each channel type, you can set up separate channels for inbound and outbound requests. This provides a simple mechanism for giving different access rights to users from outside your organization from the rights you give to users within your organization:

- To ensure that users outside your organization can only access those internal services that you choose to publish externally, you deploy those services on the inbound channel.
- To give users inside your organization access to the full range of internal and external services, you deploy those services on the outbound channel.

### Installing the gateway-supplied channels

There are different installation criteria and procedures for installing each of the gateway-supplied channels. The steps to take in each case are referenced in this topic.

- Install the Apache SOAP channel 1 application (wsgwsoap1.ear) when you install the gateway application, as described in Installing the gateway into an application server that is part of a deployment manager cell and Installing the gateway into a stand-alone application server.
- Install the Apache SOAP channel 2 application (wsgwsoap2.ear) as described in the step “Install additional gateway applications” in Installing the gateway into a deployment manager cell and Installing the gateway into a stand-alone application server.
- Install the SOAP over HTTP channel 1 application (wsgwsoaphttp1.ear) when you install the gateway application, as described in Installing the gateway into an application server that is part of a deployment manager cell and Installing the gateway into a stand-alone application server.
- Install the SOAP over HTTP channel 2 application (wsgwsoaphttp2.ear) as described in the step “Install additional gateway applications” in Installing the gateway into an application server that is part of a deployment manager cell and Installing the gateway into a stand-alone application server.
- **5.1+** Install the synchronous SOAP over Java Messaging Service (JMS) channel applications (wsgwsoapjms1.ear and wsgwsoapjms2.ear) as described in Installing the synchronous SOAP over JMS channel.

### ***Installing the synchronous SOAP over JMS channel:***

The synchronous SOAP over Java Messaging Service (JMS) channels use a JMS provider (either the JMS provider that is embedded in WebSphere Application Server, or another provider such as WebSphere MQ) to pass SOAP messages over JMS. To install and configure the JMS provider, then install the synchronous SOAP over JMS channels, complete the following steps:

1. Install a JMS provider. You might have already installed a JMS provider when you installed WebSphere Application Server. If not, install one now as described in Installing and configuring a JMS provider.
2. Use the administrative console to create and configure queue connection factories and queue destinations as described in Configuring JMS provider resources. Create two connection factories (one for each channel) and two destinations. The JMS resources and JNDI names that the synchronous SOAP over JMS channels expect by default are provided in the following table. If you use different resources and names in this step, then change the defaults when you subsequently install the channel EAR files.

<b>JMS resource</b>	<b>default JNDI name (channel 1)</b>	<b>default JNDI name (channel 2)</b>
connection factory	jms/SOAPJMSFactory1	jms/SOAPJMSFactory2
destination	jms/SOAPJMSQueue1	jms/SOAPJMSQueue2

3. Optional: If you are using embedded messaging, use the administrative console to add the two new destinations to the list of JMS Server destination names for your application server as described in Managing WebSphere internal JMS servers. Ensure that the Initial State is started.
4. Use the administrative console to create and configure listener ports in the message listener service for your application server as described in Adding a new listener port. Create two listener ports (one for each channel). The default JMS resources and associated names that the synchronous SOAP over JMS channels expect are provided in the following table. If you use different resources and names in this step, then change the defaults when you subsequently install the channel EAR files.

<b>JMS resource</b>	<b>default name (for use with SOAP over JMS channel 1)</b>	<b>default name (for use with SOAP over JMS channel 2)</b>
listener port	SOAPJMSPort1	SOAPJMSPort2
connection factory	jms/SOAPJMSFactory1	jms/SOAPJMSFactory2
destination	jms/SOAPJMSQueue1	jms/SOAPJMSQueue2

5. Use the administrative console to create a new URL provider for your application server as described in "URL provider settings" in the online help and the Information Center. Enter the property values provided in the following table:

URL provider property	value
Name	JMSHandler
Stream Handler Class Name	com.ibm.ws.protocol.Handler
Protocol	jms

6. Install the synchronous SOAP over JMS channel applications (wsgwsoapjms1.ear and wsgwsoapjms2.ear) as described in the next to last step of Installing the gateway into an application server that is part of a deployment manager cell and Installing the gateway into a stand-alone application server.
7. Optional: If you want to change the default values for listener port and connection factory for the synchronous SOAP over JMS channel applications, modify the installed applications by completing the following steps for each channel application:
  - a. In the administrative console navigation pane, click **Applications > Enterprise Applications**. A list of all installed enterprise applications is displayed.
  - b. Click the name of the channel application (for example SOAPJMSChannel1). Configuration details for the application are displayed.
  - c. Scroll down through the configuration details to the Additional Properties pane, then click Provide Listener Ports for Messaging Beans. The Binding Enterprise Beans to Listener Port Name pane is displayed.
  - d. Change the listener port value from (for example) SOAPJMSPort1 to the value that you configured in the message listener service in a previous step, then click **OK** to return to the Additional Properties pane.
  - e. In the Additional Properties pane, click Map resource env entry references to resources. The Mapping Resource Environment References to Resources pane is displayed.
  - f. Change the connection factory value from (for example) jms/SOAPJMSFactory1 to the value that you configured in the message listener service in a previous step, then click **OK** to return to the Additional Properties pane.
  - g. Scroll up through the configuration details to the Messages pane, then click **Save**. The Save to Master Configuration pane is displayed.
  - h. In the Save to Master Configuration pane, click **Save**. The changes have been committed and installed.
  - i. Stop then restart the application server.

You are now ready to Deploy the synchronous SOAP over JMS channels to the gateway.

**Web services gateway - Channel deployment details:** Here are the deployment details for the channels supplied with the Web services gateway:

- **Apache SOAP channel 1**

**Note:** The Apache SOAP channel is deprecated and wholly superseded by the SOAP over HTTP channel.

- **Channel Name:** ApacheSOAPChannel1
- **Home Location:** ApacheSOAPChannel1Bean
- **End Point Address:**
  - **Application server not within an unlocked configuration deployment:**  
http://domain\_name/wsgwsoap1
  - **Application server within an unlocked configuration deployment:**  
http://domain\_name/wsgwsoap1, http://host\_name:port\_number/wsgwsoap1

where *host\_name* and *port\_number* are the host name and port number for this application server.

- **Async Reply Context Name:** Leave blank. This function is not supported by this channel.
- **Async Reply Context Value:** Leave blank. This function is not supported by this channel.

- **Apache SOAP channel 2**

**Note:** The Apache SOAP channel is deprecated and wholly superseded by the SOAP over HTTP channel.

- **Channel Name:** ApacheSOAPChannel2
- **Home Location:** ApacheSOAPChannel2Bean
- **End Point Address:**

- **Application server not within an unlocked configuration deployment:**

- `http://domain_name/wsgwsoap2`

- **Application server within an unlocked configuration deployment:**

- `http://domain_name/wsgwsoap2, http://host_name:port_number/wsgwsoap2`

- where *host\_name* and *port\_number* are the host name and port number for this application server.

- **Async Reply Context Name:** Leave blank. This function is not supported by this channel.
- **Async Reply Context Value:** Leave blank. This function is not supported by this channel.

- **SOAP over HTTP channel 1**

- **Channel Name:** SOAPHTTPChannel1
- **Home Location:** SOAPHTTPChannel1Bean

- **5.1 End Point Address:**

- **Application server not within an unlocked configuration deployment:**

- `http://domain_name/wsgwsoaphttp1`

- **Application server within an unlocked configuration deployment:**

- `http://domain_name/wsgwsoaphttp1, http://host_name:port_number/wsgwsoaphttp1`

- where *host\_name* and *port\_number* are the host name and port number for this application server.

- **Async Reply Context Name:** Leave blank. This function is not supported by this channel.
- **Async Reply Context Value:** Leave blank. This function is not supported by this channel.

- **SOAP over HTTP channel 2**

- **Channel Name:** SOAPHTTPChannel2
- **Home Location:** SOAPHTTPChannel2Bean

- **5.1 End Point Address:**

- **Application server not within an unlocked configuration deployment:**

- `http://domain_name/wsgwsoaphttp2`

- **Application server within an unlocked configuration deployment:**

- `http://domain_name/wsgwsoaphttp2, http://host_name:port_number/wsgwsoaphttp2`

- where *host\_name* and *port\_number* are the host name and port number for this application server.

- **Async Reply Context Name:** Leave blank. This function is not supported by this channel.
- **Async Reply Context Value:** Leave blank. This function is not supported by this channel.

- **5.1+ Synchronous SOAP over Java Messaging Service (JMS) channel 1**

- **Channel Name:** SOAPJMSChannel1
- **Home Location:** SOAPJMSChannel1Bean
- **End Point Address:**

- Specify the properties of the synchronous SOAP over JMS channel 1 endpoint, using the following syntax:

- `jms:/queue or topic indicator?property_name=property_value` and so on, separating each property using the & character.

- For example, if you use the default values for queue destination and queue connection factory when you install the synchronous SOAP over JMS channels, then the first part of the end point address is:

- `jms:/queue?destination=jms/SOAPJMSQueue1&connectionFactory=jms/SOAPJMSFactory1`

- For each of the synchronous SOAP over JMS channels, here is the full list of properties that can be specified in the end point address:

Property name	Property description
<b>Destination-related Properties (required)</b>	
connectionFactory	The JNDI name of the queue or topic connection factory.
destination	The JNDI name of the destination queue or topic.
<b>JNDI-related Properties (optional)</b>	
initialContextFactory	The name of the initial context factory to use (this is mapped to the <code>java.naming.factory.initial</code> property).
jndiProviderURL	The JNDI provider Web address (this is mapped to the <code>java.naming.provider.url</code> property).
<b>JMS-related Properties (optional)</b>	
deliveryMode	An indication as to whether or not the request message is persistent. The valid values are 1 (non persistent) and 2 (persistent). The default value is 1.
timeToLive	The lifetime (in milliseconds) of the request message. A value of 0 indicates an infinite lifetime.
priority	The JMS priority associated with the request message. Valid values are 0 - 9. The default value is 4.
userid	The User ID that is used to gain access to the connection factory.
password	The password that is used to gain access to the connection factory.

- **Async Reply Context Name:** Leave blank. This function is not supported by this channel.
- **Async Reply Context Value:** Leave blank. This function is not supported by this channel.

- **5.1+ Synchronous SOAP over JMS channel 2**

- **Channel Name:** SOAPJMSChannel2
- **Home Location:** SOAPJMSChannel2Bean
- **End Point Address:**

Specify the properties of the synchronous SOAP over JMS channel 2 endpoint, using the following syntax:

`jms:/queue or topic indicator?property_name=property_value` and so on, separating each property using the & character.

For example, if you use the default values for queue destination and queue connection factory when you install the synchronous SOAP over JMS channels, then the first part of the end point address is:

`jms:/queue?destination=jms/SOAPJMSQueue2&connectionFactory=jms/SOAPJMSFactory2`

For the full list of properties that can be specified in the end point address for synchronous SOAP over JMS channel 2, see the list of properties detailed above for Synchronous SOAP over JMS channel 1.

- **Async Reply Context Name:** Leave blank. This function is not supported by this channel.
- **Async Reply Context Value:** Leave blank. This function is not supported by this channel.

## Listing and managing gateway-deployed channels

To list the channels that are currently deployed to the Web services gateway, and to view and modify their deployment details, complete the following steps:

1. Display the Web services gateway administrative user interface.
2. In the navigation pane, click the following link:

### Channels

- [List](#)

The main pane is updated with a list of all the channels that are deployed to the Web services gateway.

3. Click the name of a channel in the list. A form is displayed that you can use to view and modify the current deployment details for this channel.

4. Modify the following deployment details:

**Home Location**

Type the name of the new home for this channel.

**End Point Address**

Type the new address on which the channel is to listen.

**5.1** For an Apache SOAP or SOAP over HTTP channel running on an application server that is not part of an unlocked configuration deployment, this address is used to access the channel through the associated HTTP server. For either of these types of channel running on an application server that is part of an unlocked configuration deployment, this address is in two parts (separated by a comma):

- The address for accessing the channel through the associated HTTP server.
- The address for accessing the channel through the particular application server on which the channel application is running.

5. If the channel supports asynchronous messaging, type the values provided in the deployment documentation for the channel into the following two fields. Otherwise leave them blank.

**Async Reply Context Name**

**Async Reply Context Value**

6. To start this channel, click **YES**. To stop this channel, click **NO**.

7. Click **Apply changes**.

If the processing completes successfully, the list of deployed channels is redisplayed. Otherwise, an error message is displayed.

## Deploying channels to the Web services gateway

Before you can deploy a channel, you must install the channel application in WebSphere Application Server as described in Installing the gateway-supplied channels.

If you want to deploy the channels supplied with the Web services gateway, see the deployment details listed in Web services gateway - channel deployment details.

To deploy a channel, complete the following steps:

1. Display the Web services gateway administrative user interface.
2. In the navigation pane, click the following link:

**Channels**

- [Deploy](#)

A form is displayed that you can use to specify the deployment details.

3. Type the following channel deployment details:

**Channel Name**

Type the name by which the channel is known within the Web services gateway and by which it is listed using the **Channels > List** option. This name must be unique within the gateway.

**Home Location**

Type the name of the home for this channel.

**End Point Address**

Type the address on which the channel is to listen.

**5.1** For an Apache SOAP or SOAP over HTTP channel running on an application server that is not part of an unlocked configuration deployment, this address is used to access the



channel through the associated HTTP server. For either of these types of channel running on an application server that is part of an unlocked configuration deployment, this address is in two parts (separated by a comma):

- The address for accessing the channel through the associated HTTP server.
  - The address for accessing the channel through the particular application server on which the channel application is running.
4. If the channel supports asynchronous messaging, type the values provided in the deployment documentation for the channel into the following two fields. Otherwise leave them blank.  
**Async Reply Context Name**  
**Async Reply Context Value**
  5. Click **OK**.

If the processing completes successfully, the list of deployed channels is updated to include the new channel. Otherwise, an error message is displayed.

To use this channel, select it for use with a Web service as described in *Deploying Web services to the Web services gateway and Listing and managing gateway-deployed Web services*.

## Removing channels from the Web services gateway

To remove a channel, complete the following steps:

1. Display the Web Services gateway administrative user interface.
2. In the navigation pane, click the following link:

**Channels**  
• [Remove](#)

The main pane is updated with a list of all the channels that are deployed to the Web services gateway. Next to each entry in the list is a check box, and information on the number of Web services that currently use the channel.

3. Click the name of a channel in the list. A form is displayed that you can use to view the current deployment details for this channel, including a list of the Web services that currently use the channel.
4. Select the check box for every channel that you want to remove.

**Note:** When you remove a channel that is currently used by one or more Web services, the gateway removes the channel from the channel list for each associated Web service.

5. Click **OK**.

If the processing completes successfully, the list of deployed channels is updated. Otherwise, an error message is displayed.

## Working with filters

**Note:** **5.1+** The use of filters is deprecated. However, you can achieve similar results using Java API for XML-based Remote Procedure Call (JAX-RPC) handlers. For more information, see *JAX-RPC handlers - An alternative to gateway filters*.

Before you can work with a filter, you must install the filter application in WebSphere Application Server as described in the next to last step of *Installing the gateway into a deployment manager cell* and *Installing the gateway into a stand-alone application server*.

You should not make changes to the filter deployment settings for a gateway that is part of an unlocked configuration load-balancing group unless you are making changes, under controlled conditions, to all members of the group.



From the navigation pane of the Web services gateway administrative user interface, you can choose the following actions for **Filters**:

- **List** to list the deployed filters, and modify their deployment details.
- **Deploy** to deploy a filter.
- **Remove** to remove filters.

You can also write your own filters for the Web services gateway.

## Filters - Service interceptors for the Web services gateway

Filters are used to intercept service invocations and responses that come into and leave the gateway.

**Note:** **5.1+** The use of filters is deprecated. However, you can achieve similar results using Java API for XML-based Remote Procedure Call (JAX-RPC) handlers. For more information, see JAX-RPC handlers - An alternative to gateway filters.

Filters can perform a wide range of tasks, from logging messages, to transforming the content, to terminating an incoming request. Filters are deployed to the Web services gateway as described in Deploying filters to the Web services gateway, then registered for use with individual Web services as described in Working with Web services.

## Listing and managing gateway-deployed filters

**Note:** **5.1+** The use of filters is deprecated. However, you can achieve similar results using Java API for XML-based Remote Procedure Call (JAX-RPC) handlers. For more information, see JAX-RPC handlers - An alternative to gateway filters.

To list the filters that are currently deployed to the Web services gateway, and to view and modify their deployment details, complete the following steps:

1. Display the Web Services gateway administrative user interface.
2. In the navigation pane, click the following link:

### Filters

- [List](#)

The main pane is updated with a list of all the filters that are deployed to the Web services gateway.

3. Click the name of a filter in the list. A form is displayed that you can use to view and modify the current deployment details for this filter.

4. Modify the following deployment detail:

### Home Location

Type the name of the new home for this filter.

5. Click **Apply changes**.

If the processing completes successfully, the list of deployed filters is redisplayed. Otherwise, an error message is displayed.

## Deploying filters to the Web services gateway

Use this task to deploy a filter to the Web services gateway.

**Note:** **5.1+** The use of filters is deprecated. However, you can achieve similar results using Java API for XML-based Remote Procedure Call (JAX-RPC) handlers. For more information, see JAX-RPC handlers - An alternative to gateway filters.

Before you can deploy a filter, you must install the filter application in WebSphere Application Server as described in the next to last step of Installing the gateway into an application server that is part of a deployment manager cell and Installing the gateway into a stand-alone application server.

**Note:** You can deploy multiple instances of a filter by entering different filter names.

To deploy a filter, complete the following steps:

1. Display the Web services gateway administrative user interface.
2. In the navigation pane, click the following link:

**Filters**

- [Deploy](#)

A form is displayed for you to specify the deployment details.

3. Type the following filter deployment details:

**Filter Name**

Type the name by which the filter is known within the Web services gateway, and by which it is listed using the **Filters > List** option. This name must be unique within the Web services gateway.

**Home Location**

Type the JNDI name for this filter.

4. Click **OK**.

If the processing completes successfully, the list of deployed filters is updated to include the new filter. Otherwise, an error message is displayed.

To use this filter, select it for use with a Web service as described in [Deploying Web services to the Web services gateway](#) and [Listing and managing gateway-deployed Web services](#).

## Removing filters from the Web services gateway

**Note:** **5.1+** The use of filters is deprecated. However, you can achieve similar results using Java API for XML-based Remote Procedure Call (JAX-RPC) handlers. For more information, see [JAX-RPC handlers - An alternative to gateway filters](#).

To remove a filter, complete the following steps:

1. Display the Web Services gateway administrative user interface.
2. In the navigation pane, click the following link:

**Filters**

- [Remove](#)

The main pane is updated with a list of all the filters that are deployed to the Web services gateway. Next to each entry in the list is a check box, and information on the number of Web services that currently use the filter.

3. Click the name of a filter in the list. A form is displayed that you can use to view the current deployment details for this filter, including a list of the Web services that currently use the filter.
4. Select the check box for every filter that you want to remove.

**Note:** When you remove a filter that is currently used by one or more Web services, the gateway removes the filter from the filter lists for each associated Web service.

5. Click **OK**.

If the processing completes successfully, the list of deployed filters is updated. Otherwise, an error message is displayed.

## Working with JAX-RPC handlers

For information about creating JAX-RPC handlers, see the IBM developerWorks article [A developer introduction to JAX-RPC, Part 2: Mine the JAX-RPC specification to improve Web service interoperability](#).

Before you can work with a Java API for XML-based Remote Procedure Call (JAX-RPC) handler, you must make the handler classes available to the application server in one of the following ways:

- Copy the individual class file into a directory structure under *install\_root/classes* that matches the package name of the class, where *install\_root* is the root directory for the installation of your application server (by default WebSphere/AppServer). For example a handler class `com.ibm.jaxrpc.handler.TestHandler` is copied into *install\_root/classes/com/ibm/jaxrpc/handler*.
- Package the class files for all your handlers as a JAR file, then copy it into *install\_root/lib/app*.

You should not make changes to the handler deployment settings for a gateway that is part of an unlocked configuration load-balancing group unless you are making changes, under controlled conditions, to all members of the group.

From the navigation pane of the Web services gateway administrative user interface, you can choose the following actions for **Handlers**:

- **List** to list the deployed handlers, and modify their deployment details.
- **Deploy** to deploy a handler.
- **Remove** to remove handlers.

### JAX-RPC handlers - An alternative to gateway filters

JAX-RPC handlers interact with messages as they pass between the service requester and the gateway, and between the gateway and the target service.

A Java API for XML-based remote procedure call (JAX-RPC) handler is a Java class that performs a similar range of tasks to a gateway filter - from logging messages, to transforming their contents, to terminating an incoming request. For more detailed information about JAX-RPC handlers, see the IBM developerWorks article [A developer introduction to JAX-RPC, Part 2: Mine the JAX-RPC specification to improve Web service interoperability](#).

All new gateway installations should use JAX-RPC handlers rather than gateway filters, for the following reasons:

- JAX-RPC is part of the proposed Java 2 Platform, Enterprise Edition (J2EE) 1.4, and JAX-RPC handlers are emerging as the standard approach in Java for intercepting and filtering service messages.
- JAX-RPC handlers are already being widely implemented - and any JAX-RPC handlers you write for use in other systems can also be deployed to the gateway.
- JAX-RPC handlers are already accepted as the standard approach in Java for managing message-level security as defined by the Web Services Security (WS-Security) specification.

Filters are applied at the level of the gateway service. JAX-RPC handlers are applied to:

- The gateway service and the channel (for messages passing between the service requester and the gateway).
- The target service and the target service port (for messages passing between the gateway and the target service).

JAX-RPC handlers are deployed to the Web services gateway as described in [Deploying handlers to the Web services gateway](#) and [Listing and managing gateway-deployed handlers](#). After you deploy a handler to the gateway you can add it to the handler chain for the gateway service and the channel, or to the handler chain for the target service and the target service port, as described in [Listing and managing gateway-deployed Web services](#).

If you deploy a gateway service to run in proxy mode, then you must also deploy (to the gateway service and the channel) a JAX-RPC handler that can set the target endpoint for the gateway service. For more information, see [JAX-RPC handlers and proxy operation](#).

Use of JAX-RPC handlers is currently supported only by the SOAP over HTTP channel.

## JAX-RPC handlers and proxy operation

You can set the gateway to act purely as a proxy for your service, then use JAX-RPC handlers to set the endpoints for incoming request messages for the service.

When you deploy a service to the gateway and enable proxy operation, the gateway takes no action with regard to that service other than to invoke it as a WSIF service. When you deploy a service to run in proxy mode, you also enable selective SOAP parsing for that service, and deploy a JAX-RPC handler that sets the `transport.url` context property to set the target endpoint for the service.

When the gateway receives a message, it needs to know whether the request being invoked is request and response, or one-way. Because the Gateway is not parsing the SOAP message, it cannot get this information from the message. Therefore the requesting clients append an `operationMode` HTTP query string parameter to the Web address for the gateway service. The value of this parameter is either `oneway` or `requestResponse`. For example, if the Web address of the deployed proxy service is `http://.../wsgwsoap/urn:wsgw::ProxyService`, then requesting clients indicate that they are sending a one-way request by using the following URL:

```
http://.../wsgwsoap/urn:wsgw::ProxyService?operationMode=oneway
```

If the `operationMode` parameter is missing, the gateway assumes that the requested method is `requestResponse`.

A gateway service that runs in proxy mode has no gateway-deployed target services and therefore no WSDL that the gateway can use to configure the WSIF service invocation. A generic template WSDL file is used to configure the basic parameters for the invocation call (for example which WSIF provider to use), but you can override the default by supplying your own equivalent template. The default template WSDL defines a single `portType` with two operations: `oneway` and `requestResponse`. If the operation mode is `oneway`, then the gateway selects the one-way operation from the WSDL. The default template WSDL is located at `http://.../wsgw/admin/ProxyServiceTemplate.wsdl`.

If you want the gateway to use a different template WSDL file, then you specify the Web address of the new template WSDL file when you deploy the service to the Web services gateway. Your new template WSDL must implement the same port type, binding, service, and port names as the default template. The only differences that can exist are in the extension elements used to configure the binding. In your new template WSDL, the value of the `<soap:target address>` tag must be a properly formatted Web address but it does not have to point to a real page. For example, a value of `this.is.a.fake.url` is rejected, whereas a value of `http://this.is.a.fake.url` is accepted. The JAX-RPC handler that sets the `transport.url` property overrides this value at runtime with the real Web address.

**Note:** If the JAX-RPC handler is not deployed, then the gateway attempts to send all requests to the fake url specified in the `<soap:target address>` tag in the template WSDL.

Proxy operation is currently supported only by the SOAP over HTTP channel.

For the detailed steps to take to enable proxy operation for a Web service, see [Deploying a gateway service for proxy operation](#).

## Listing and managing gateway-deployed handlers

To list the handlers that are currently deployed to the Web services gateway, and to view and modify their deployment details, complete the following steps:

1. Display the Web Services gateway administrative user interface.
2. In the navigation pane, click the following link:

### Handlers

- [List](#)

The main pane is updated with a list of all the handlers that are deployed to the Web services gateway.

3. Click the name of a handler in the list. A form is displayed that you can use to view and modify the current deployment details for this handler.
4. Modify the following deployment details. Click the **Add** and **Remove** buttons to the right of each row to add and remove corresponding parameters of the handler (the **Remove** button appears automatically after you add your first parameter).

**Handler Class**

Modify the name of the class that is to be instantiated.

**Init Parameter name/value**

Add or remove up to three initial parameter names and associated values.

**SOAP Header QName**

Add or remove the QName of a SOAP header that will be processed by the handler.

**SOAP role**

Add or remove a SOAP actor definition that the Handler will play as a role. For more information, see the SOAP specification.

5. Click **Apply changes**.

If the processing completes successfully, the list of deployed handlers is redisplayed. Otherwise, an error message is displayed.

## Deploying handlers to the Web services gateway

Use this task to deploy a Java API for XML-based Remote Procedure Call (JAX-RPC) handler to the gateway.

For information about creating JAX-RPC handlers, see the IBM developerWorks article *A developer introduction to JAX-RPC, Part 2: Mine the JAX-RPC specification to improve Web service interoperability*.

Before you can deploy a handler, you must make the handler class available to the application server in one of the following ways:

- Copy the individual class file into a directory structure under *install\_root/classes* that matches the package name of the class, where *install\_root* is the root directory for the installation of your application server (by default WebSphere/AppServer). For example a handler class `com.ibm.jaxrpc.handler.TestHandler` is copied into *install\_root/classes/com/ibm/jaxrpc/handler*.
- Package the class files for all your handlers as a JAR file, then copy it into *install\_root/lib/app*.

**Note:** You can deploy multiple instances of a handler by entering different handler names.

To deploy a handler, complete the following steps:

1. Display the Web services gateway administrative user interface.
2. In the navigation pane, click the following link:

**Handlers**

- [Deploy](#)

A form is displayed for you to specify the deployment details.

3. Type the following required handler deployment details:

**Handler Name**

Type the name by which the handler is known within the Web services gateway, and by which it is listed using the **Handlers > List** option. This name must be unique within the Web services gateway. For example `TestHandler`.

**Handler Class**

Type the name of the class that is to be instantiated. For example `com.ibm.jaxrpc.handler.TestHandler`.

4. Optional: Type the following optional handler deployment details:

**Init Parameter name/value**

Type up to three initial parameter names and associated values.

**SOAP Header QName**

Type the QName of a SOAP header that will be processed by the handler.

**SOAP role**

Type a SOAP actor definition that the Handler will play as a role. For more information, see the SOAP specification.

5. Click **OK**.

If the processing completes successfully, the list of deployed handlers is updated to include the new handler. Otherwise, an error message is displayed.

To use this handler, select it for use with a Web service as described in *Deploying Web services to the Web services gateway and Listing and managing gateway-deployed Web services*.

## Removing handlers from the Web services gateway

Use this task to remove a Java API for XML-based Remote Procedure Call (JAX-RPC) handler from the gateway.

To remove a handler, complete the following steps:

1. Display the Web Services gateway administrative user interface.
2. In the navigation pane, click the following link:

**Handlers**

- [Remove](#)

The main pane is updated with a list of all the handlers that are deployed to the Web services gateway. Next to each entry in the list is a check box, and information on the number of Web services that currently use the handler.

3. Click the name of a handler in the list. A form is displayed that you can use to view the current deployment details for this handler, including a list of the Web services that currently use the handler.
4. Select the check box for every handler that you want to remove.

**Note:** When you remove a handler that is currently used by one or more Web services, the gateway removes the handler from the handler lists for each associated Web service.

5. Click **OK**.

If the processing completes successfully, the list of deployed handlers is updated. Otherwise, an error message is displayed.

## Working with UDDI references

You should not make changes to the UDDI references for a gateway that is part of an unlocked configuration load-balancing group.

A UDDI reference is a pointer to a UDDI registry. This registry can be a private UDDI registry such as the IBM WebSphere UDDI Registry, or a public UDDI registry.

In the UDDI model, Web services are owned by businesses, and businesses are owned by Authorized Names. Each UDDI reference gives access to the Web services that are owned by a single Authorized Name in a single UDDI registry.

From the navigation pane of the Web services gateway administrative user interface, you can choose the following actions for **UDDI References**:

- **List** to list the deployed UDDI references, and modify their deployment details.



- **Deploy** to deploy a UDDI reference.
- **Remove** to remove UDDI references.

For more information about how the gateway works with UDDI registries, see [UDDI registries - Web service directories that integrate with the Web services gateway](#) and [Publishing a Web service to a UDDI registry for deployment to the gateway](#). For more general information about UDDI and UDDI registries, see the UDDI community at [uddi.org](http://uddi.org).

## **UDDI registries - Web service directories that integrate with the Web services gateway**

### **UDDI**

The Universal Description, Discovery and Integration (UDDI) specification defines a way to publish and discover information about Web services.

In this specification:

- Each Web service is owned by one business, and each business (and the Web services it owns) is maintained by one Authorized Name.
- One Authorized Name can own many businesses, and one business can own many Web services.

The UDDI specification also associates Web services with Technical models. Using these models or generic categories, a UDDI registry user can search for a type of service, rather than needing to know the access details for a specific service.

For more general information about UDDI, see the UDDI community at [uddi.org](http://uddi.org).

### **UDDI registries**

UDDI registries use the UDDI specification to publish directory listings of Web services. There are Universal Business Registries (sometimes referred to as public UDDI registries) hosted worldwide, including one hosted by IBM. Enterprises can also host their own internal registries behind their firewalls (sometimes referred to as private UDDI registries) to better manage their internal implementation of Web services. The IBM WebSphere UDDI Registry is an example of a private UDDI registry.

### **How the gateway interacts with UDDI registries**

The gateway interacts with UDDI registries in two ways:

- When you deploy a Web service to the gateway, you specify the location of the internal WSDL file that describes the Web service to be deployed. This WSDL file can be located through a UDDI registry.
- For any gateway-deployed Web service, you can tell the gateway to create entries for the Web service in one or more UDDI registries.

To enable your gateway to interact with a UDDI registry, you create one or more gateway pointers to the registry. The gateway refers to these pointers as UDDI references, and you create them as described in [Deploying UDDI references to the Web services gateway](#). Each UDDI reference includes the following parameters:

- The access points for the UDDI registry (the **Inquiry URL** and the **Publish URL**).
- The Authorized Name (the **User ID** and **Password**) for the owner of one or more businesses in the UDDI registry.

You get the Authorized Name from the target UDDI registry. For more information, see [Publishing a Web service to a UDDI registry for deployment to the gateway](#).



A given UDDI reference can only access the Web services that are owned by the businesses that are owned by a single Authorized Name. Therefore if you need to access two Web services in the same registry, and each service is owned by a different “Authorized Name”, then you need to create two UDDI references.

When you deploy a Web service, and you specify that the internal WSDL file is located through a UDDI registry, you enter the following two parameters:

- The UDDI reference that can access this service.
- The service key that the UDDI registry has assigned to this service.

You get the service key from the target UDDI registry. For more information, see Publishing a Web service to a UDDI registry for deployment to the gateway.

When you tell the gateway to create entries for a deployed Web service in one or more UDDI registries, you enter the following two parameters:

- The UDDI references (one for each registry) that can access the UDDI business category under which you want to publish this service.
- The business key that identifies the UDDI business category.

You get the business key from the target UDDI registry. For more information see Publishing a Web service to a UDDI registry for deployment to the gateway.

Because the gateway only interacts with UDDI registries at the level of specific Web services, the gateway does not use UDDI Technical models.

## Listing and managing gateway-deployed UDDI references

To list the UDDI references that are currently deployed to the Web services gateway, and to view and modify their deployment details, complete the following steps:

1. Display the Web services gateway administrative user interface.
2. In the navigation pane, click the following link:

### UDDI References

- [List](#)

The main pane is updated with a list of all the UDDI references that are deployed to the Web services gateway.

3. Click the name of a UDDI reference in the list. A form is displayed that you can use to view and modify the current deployment details for this UDDI reference.
4. Modify the following deployment details:

#### **Inquiry URL**

Type the new Web address that provides access to this registry for the SOAP inquiry API.

#### **Publish URL**

Type the new Web address that provides access to this registry for the SOAP publish API.

#### **User Name**

Type the new user ID for the Authorized Name that has update access to this registry.

#### **Password**

Type the password for this new user ID.

#### **Confirm Password**

Type again the password for this new user ID.

#### **Note:**

The values you enter here for **User Name** and **Password** must match those of the owner of the corresponding business in UDDI. You can see the owning user ID in UDDI by looking at the business details under the “Authorized Name” field.

If the values you enter here do not match the “Authorized Name” values for the business that owns the service, then the service will not be published or found.

If the business that owns the service has more than one Authorized Name, you might want to set up multiple UDDI references (each with a different user ID) to the same UDDI registry .

5. Click **Apply changes**.

If the processing completes successfully, the list of deployed UDDI references is redisplayed. Otherwise, an error message is displayed.

## Deploying UDDI references to the Web services gateway

To deploy a UDDI reference, complete the following steps:

1. Display the Web services gateway administrative user interface.
2. In the navigation pane, click the following link:

### UDDI References

- [Deploy](#)

A form is displayed for you to specify the deployment details.

3. Type the following UDDI reference deployment details:

#### Reference Name

Type the name by which the UDDI reference is known within the Web services gateway, and by which it is listed using the **UDDI References > List** option. This name must be unique within the gateway. You might need more than one UDDI reference for a given UDDI registry. For more information, see UDDI registries - Web service directories that integrate with the Web services gateway.

#### Inquiry URL

Type the Web address that provides access to this registry for the SOAP inquiry API.

#### Publish URL

Type the Web address that provides access to this registry for the SOAP publish API.

#### User Name

Type the user ID for an “Authorized Name” that has update access to this registry.

#### Password

Type the password for this user ID.

#### Confirm Password

Type again the password for this user ID.

#### Note:

The values you enter here for **User Name** and **Password** must match those of the Authorized Name in the UDDI registry. You can see the owning user ID in UDDI by looking at the business details under the Authorized Name field.

4. Click **OK**.

If the processing completes successfully, the list of deployed UDDI references is updated to include the new UDDI reference. Otherwise, an error message is displayed.

To use this UDDI reference, select it for use with a Web service as described in *Deploying Web services to the Web services gateway and Listing and managing gateway-deployed Web services*.

## Removing UDDI references from the Web services gateway

To remove a UDDI reference, complete the following steps:

1. Display the Web services gateway administrative user interface.
2. In the navigation pane, click the following link:

## UDDI References

- [Remove](#)

The main pane is updated with a list of all the UDDI references that are deployed to the Web services gateway. Next to each entry in the list is a check box, and information on the number of Web services that currently use the UDDI reference.

3. Click the name of a UDDI reference in the list. A form is displayed that you can use to view the current deployment details for this UDDI reference, including a list of the Web services that currently use the UDDI reference.
4. Select the check box for every UDDI reference that you want to remove. When you remove a UDDI reference that is currently used by one or more Web services, the gateway removes the UDDI reference from the UDDI reference list for each associated Web service.
5. Click **OK**.

If the processing completes successfully, the list of deployed UDDI references is updated. Otherwise, an error message is displayed.

## Working with Web services

If you change the namespace URI, you break the link back to the Web services gateway for every Web service that you have already deployed. Therefore you must set the namespace URI before you deploy any Web services.

When you configure a Web service, you associate it with one or more of the following resources:

- The channels on which the service is available.
- Any filters that apply to the service.
- **5.1 +** Any JAX-RPC handlers that apply to the service.
- Any security bindings that apply to the service.
- Any references to UDDI registries in which entries for the service are created.

You choose each of these resources from a list of gateway-deployed resources, so you should (if possible) deploy each resource to the gateway before you deploy a Web service that uses it.

You should not make changes to the Web service deployment settings for a gateway that is part of an unlocked configuration load-balancing group unless you are making changes, under controlled conditions, to all members of the group.

From the navigation pane of the Web services gateway administrative user interface, choose one of the following actions for **Services**:

- **List** to list the deployed Web services, and modify their deployment details.
- **Deploy** to deploy a Web service.
- **Remove** to remove Web services.

## Target services and gateway services

The gateway makes the following distinction between a target service and a gateway service:

- A target service is an actual service that exists outside the gateway, and that you make available to others through the gateway.
- A gateway service is the view of a target service that the gateway gives to service requesters. It is decoupled from the actual target service.

When you deploy a target service to the gateway, the gateway creates an equivalent gateway service. This gateway service is described in a new representation of the target service WSDL that is published to a gateway-controlled URL. This indirection gives the following benefits:

- You can move the target service to a new location, or replace it with a new implementation, and you only need to update the target service information that is held in the gateway. Existing service requesters can still find it and use it, because (as far as they can see) nothing has changed.
- If you have several different implementations of the same service, and you deploy them all to the gateway as multiple target services for a single gateway service, then they appear to service requesters as a single service. You can then use a filter (or similar mechanism) to choose the most appropriate target service for each incoming request.
- You can set, quite independently, the security measures that apply between the service requester and the gateway, and the security measures that apply between the gateway and each target service.

## Listing and managing gateway-deployed Web services

Use this task to list the Web services that are deployed to the Web services gateway, and modify their deployment details.

Do not deploy multiple target services to the same gateway service unless you have a filter (or similar mechanism) that can select the required target service, because without filtering the gateway always picks the first target service in the list.

For information on how to write pluggable filters that can select the target service, see [Writing a filter for the Web services gateway](#) and [Using a filter to select a target service and port](#).

To list the Web services that are currently deployed to the Web services gateway, and to view and modify their deployment details (including adding or removing multiple target services) complete the following steps:

1. Display the Web services gateway administrative user interface.
2. In the navigation pane, click the following link:

### Services

- [List](#)

The main pane is updated with a list of all the Web services that are deployed to the Web services gateway.

3. Click the name of a Web service in the list. A form is displayed that you can use to view and modify the current deployment details for this gateway service, and to add or remove multiple target services.
4. In the **Gateway Service Properties** section you can change the following details. When you have finished making changes, click **Apply Changes**.
  - a. **Authorization Policy - Control access to this service.** Use this check box to enable or disable operation-level authorization for this gateway service.
  - b. **Audit Policy - Log requests to this service.**  
 The Audit policy indicates whether the MessageWarehouse object, if present, is used to log requests and responses for this service. If you have a Message Warehouse implementation, use this check box to enable or disable logging of requests and responses for this Web service.  
 For more information, see [Capturing Web service invocation information](#) and [The MessageWarehouse interface](#).
  - c. In this release of the gateway, the **Annotation URL** field is not used.
  - d. If you want to publish the service to one or more UDDI registries (selected in the subsequent **UDDI References** section), enter the UDDI business key in the field provided under **UDDI Publication Properties**. This key identifies the business category under which you want your service to appear in UDDI. To get a list of valid business keys, look up businesses in a UDDI registry. Here is an example of a UDDI business key: 08A536DC-3482-4E18-BFEC-2E2A23630526.  
 For more information about UDDI business keys, see [Publishing a Web service to a UDDI registry for deployment to the gateway](#).
  - e. If you want to enable or disable **Service Security**, select or clear the **Enable service security** check box. If you choose to enable service security, then you must use the **Edit service security**

**configuration** option to configure secure communication for this gateway service between the service requester (the client) and the gateway. For more information, see Editing the service security configuration.

5. In the **Target Services** section you add, modify or remove services from a list of target services for this single gateway service. Every target service on this list provides exactly the same service, and they are presented by the gateway to the service requesters as a single gateway service. To add a new target service, complete the following steps:
  - a. **WSDL Location and Location Type.** Specify the location, and select the location type, for the internal WSDL file that describes the Web service to be deployed. The WSDL file is either located at a Web address, or through a UDDI registry.

**Note:**

- When the gateway deploys the Web service, it generates a matching external WSDL file that it makes available to gateway users. This external WSDL file also describes the service, but is located at a new Web address and is generated and maintained by the Web services gateway.
- If the service provider makes changes to the internal WSDL file, the gateway does not pick up the changes until the application server on which the gateway is hosted is restarted. For performance reasons, the gateway does not check the service information on each service invocation. The gateway reads the service information when it deploys the Web service, then uses this information for as long as the service remains deployed, or until the application server is restarted.
- If the Web service is also being published to one or more UDDI registries, then the internal WSDL file is required to remove the service from the gateway.

If the location for the internal WSDL file is a Web address, type the Web address. If the binding and service definition for this Web service are held in separate WSDL files, then type the Web address of the WSDL file that defines the binding.

If the WSDL is located through a UDDI registry, type *uddiReference,serviceKey* where

- *uddiReference* is the reference name by which a currently-deployed UDDI reference is known within the gateway, and by which it is listed using the **UDDI References > List** option.
- *serviceKey* is the service key that the UDDI registry has assigned to the service. Here is an example of a UDDI service key: 34280367-0ECF-46CE-B804-14C21D6D0FB1. For more information about UDDI service keys, see Publishing a Web service to a UDDI registry for deployment to the gateway.

- b. **Target Service Name.** If the Web service WSDL contains more than one service, or the WSDL is located through a UDDI registry, type the target service name from the target service WSDL.
- c. **Target Service Namespace.** If the Web service WSDL contains more than one service, or the WSDL is located through a UDDI registry, type the namespace of the target service name from the target service WSDL.
- d. **Target Service Identity Information.** Type the identity by which the target service is known within the Web services gateway. This identity need not be unique.

**Note:** If you are mapping multiple target services, and also writing a filter to select a target service, you might use the Target Service Identity Information to select a particular target service from the set.

- e. Click **add**. The target service is added to a list of target services.

To modify the deployment details for an existing target service, click on the name of the service in the list of target services. A form is displayed, containing the same fields that you filled in when you added the service, and also the following additional fields:

- **Started.** Enable this check box option if you want the gateway to offer this Web service to service requesters.
- **Enable target service security.** Enable this check box if you want the gateway to use secure communication between the gateway and the target service.

- **Edit target service security configuration.** If you choose to enable target service security, then you must use this option to configure secure communication between the gateway and the target service. For more information see Editing the target service security configuration.

**5.1 +** At the bottom of this form is the **Target Service Ports** section. This section contains a list of available ports for this target service. To edit the JAX-RPC handler configuration for a port for this target service, complete the following steps:

- a. Click the link next to the port name. The **Edit Target Service JAX-RPC handler configuration** pane is displayed.
- b. In the **Edit Target Service JAX-RPC handler configuration** pane, use the **add** and **remove** buttons, and the drop-down lists, to modify the chain of handlers for this target service and target service port.
- c. Click the target service name to return to the target service deployment details form.

When you have finished making changes to the target service deployment details, click **Apply Changes**.

6. In the **Channels** section, you can add or remove channels from the list of deployed channels through which this service is available.

**5.1 +** To edit the JAX-RPC handler configuration for a channel for this gateway service, complete the following steps:

- a. Click the link next to the channel name. The **Edit Gateway Service JAX-RPC handler configuration** pane is displayed.
- b. In the **Edit Gateway Service JAX-RPC handler configuration** pane, use the **add** and **remove** buttons, and the drop-down lists, to modify the chain of handlers for this gateway service and channel.
- c. Click the gateway service name to return to the gateway service deployment details form.

7. In the **Request filters** section of the gateway service deployment details form, you can add or remove filters from the list of deployed filters that are applied to the service request message.

**Note:** The filters are executed in the order shown. To add a filter into the list at a particular position, use the **at position** menu.

8. In the **Response filters** section, you can add or remove filters from the list of deployed filters that are applied to the service response message.

**Note:** The filters are executed in the order shown. To add a filter into the list at a particular position, use the **at position** menu.

9. In the **UDDI References** section, you can add or remove UDDI references from the list of deployed UDDI references to UDDI registries in which this service is published. If you select one or more UDDI references in this step, you must also enter a UDDI business key in the field provided under **UDDI Publication Properties** as described in a previous step. For more information about how the gateway works with UDDI registries, see UDDI registries - Web service directories that integrate with the Web services gateway.
10. In the **Exported WSDL definitions** section there are two pairs of WSDL links. Both pairs link to the external WSDL implementation definition, and the external WSDL interface definition.
  - To view the XML source of the associated external WSDL for the service, use the **View external WSDL implementation definition** links.
  - To return the WSDL parsed for use by application programs that need the WSDL definitions for the service, use the pair labelled **External WSDL implementation definition (WSDL only)**.

WSDL documents are XML documents, and two pairs of links are supplied in order to display the same external WSDL in two different document formats. The pair of links labelled **View external WSDL implementation definition** display the XML source for the WSDL, whereas the pair of links labelled **External WSDL implementation definition (WSDL only)** display the WSDL as it appears in a Web browser.



For example, in XML notation the ampersand character (“&”) is reserved for use in encoding special characters such as the quotation mark (“&quot;”). The ampersand character itself is encoded as “&amp;”. When you select either of the **View external WSDL implementation definition** links, you view the literal XML source and the ampersand character is represented as “&amp;”. When you select either of the **External WSDL implementation definition (WSDL only)** links, your Web browser parses the XML source, converts the special characters back into their normal form, and the ampersand character is represented as “&”.

**Note:** The “&amp;” encoding is required only within the WSDL XML source document. You should not use it in any endpoint address that you supply to a service client. If you get the endpoint address from the WSDL programatically, then the XML parser handles the conversion. If you manually copy and paste the address into the properties file, you should copy the address from one of the **External WSDL implementation definition (WSDL only)** links.

If there is an error generating the WSDL then a blank page is returned.

To help your service users locate the WSDL documents for services that are deployed to the Web services gateway, the gateway also supports the WS-Inspection specification. To open a Web services inspection document that contains references to the WSDL documents for all of the gateway-deployed services, you issue an HTTP GET against

`http://host:port/wsgw/wsinspection.wsil`

where *host* and *port* are the host name and port number on which your HTTP server is listening.

**Note:** When you use the Web address provided in Testing the installation to list the services that are deployed to the SOAP over HTTP channel, you see a link labeled **WSDL** next to each service in the list. Do not use these links. To get the WSDL location for your service, open a Web services inspection document.

## Deploying Web services to the Web services gateway

To deploy a Web service, complete the following steps:

1. Display the Web services gateway administrative user interface.
2. In the navigation pane, click the following link:

### Services

- [Deploy](#)

A form is displayed that you use to specify the deployment details.

3. In the **Gateway Service Properties** section, specify the following details:
  - a. **Gateway Service Name.** Type the name by which the Web service is known within the gateway, and by which it is listed using the **Services > List** option. This name must be unique within the gateway and must not contain any spaces.
  - b. **Message part representation.** Choose between **Generic classes** and **Deployed Java classes**, and (if available) choose whether or not **Selective SOAP parsing** is enabled.

You can use either generic classes or deployed Java classes to represent simple, compound and complex data types, subject to the constraints described in Data type representation - Choosing between Generic classes and Deployed Java classes. If your Web service uses complex data types, and there is no other constraint that forces you to use deployed Java classes, then choose generic classes.

**5.1+** If you enable selective SOAP parsing, then the gateway parses only the headers of incoming messages for this service. The message body is passed through unchecked. Choose this option to optimize the speed at which messages pass through the gateway. Selective SOAP parsing is only supported by the SOAP over HTTP channel. For more information, see Selective SOAP parsing.

- c. **5.1+ Operate in Proxy Mode.** Choose whether or not this gateway service runs in proxy mode. If you enable proxy operation, then you must also choose selective SOAP parsing in the previous



step, and deploy a JAX-RPC handler that sets the target endpoint for this gateway service. Proxy operation is only supported by the SOAP over HTTP channel. For more information, see JAX-RPC handlers and proxy operation and Deploying a gateway service for proxy operation.

- d. **5.1+ Proxy Service Template WSDL Override.** If you enable proxy operation in the previous step, and you want to override the default proxy service template WSDL that is supplied with the gateway, then type the location (as a URI) of the replacement template WSDL file. Otherwise, leave this field blank. For more information, see JAX-RPC handlers and proxy operation.
- e. **Authorization Policy - Control access to this service.** If you want to enable operation-level authorization for this Web service, enable this check box.
- f. **Audit Policy - Log requests to this service.**

The Audit policy indicates whether the MessageWarehouse object, if present, is used to log requests and responses for this service. If you have a Message Warehouse implementation, and you want it to log requests and responses for this Web service, enable this check box.

For more information see Capturing Web service invocation information and The MessageWarehouse interface.

- g. In this release of the gateway, the **Annotation URL** field is not used.
- h. Select the deployed resources for the Web service to use, from the following lists:

**Channels**

Select one or more deployed channels through which this service is available.

**Request filters**

Select zero or more deployed filters to apply to the request.

**Response filters**

Select zero or more deployed filters to apply to the response.

**UDDI References**

Select zero or more deployed UDDI references (one for each UDDI registry) that can access the UDDI business category under which you want to publish this service. If you select one or more UDDI references in this step, you must also enter the UDDI business key in step 3h. For more information about how the gateway works with UDDI registries, see UDDI registries - Web service directories that integrate with the Web services gateway.

4. In the **Target Service Properties** section, specify details for the first target service that is represented by this gateway service.

**Note:** **5.1+** If you are deploying a service for proxy operation, do not specify any target service properties.

**Note:** If you have several implementations of the same Web service, you can map them all to the same deployed gateway service. To add more target services to this gateway service, use the Services > List option.

- a. **WSDL Location and Location Type.** Specify the location, and select the location type, for the internal WSDL file that describes the Web service to be deployed. The WSDL file is either located at a Web address, or through a UDDI registry.

**Note:**

- When the gateway deploys the Web service, it generates a matching external WSDL file that it makes available to gateway users. This external WSDL file also describes the service, but is located at a new Web address and is generated and maintained by the Web services gateway.
- If the service provider makes changes to the internal WSDL file, the gateway does not pick up the changes until the application server on which the gateway is hosted is restarted. (For performance reasons, the gateway does not check the service information on each service invocation. The gateway reads the service information when it deploys the Web service, then uses this information for as long as the service remains deployed, or until the application server is restarted.)

- If the Web service is also being published to one or more UDDI registries, then the internal WSDL file is required to remove the service from the gateway.

If the location for the internal WSDL file is a Web address, type the Web address. If the binding and service definition for this Web service are held in separate WSDL files, then type the Web address of the WSDL file that defines the binding.

If the WSDL is located through a UDDI registry, type *uddiReference*, *serviceKey* where:

- *uddiReference* is the reference name by which a currently-deployed UDDI reference is known within the gateway, and by which it is listed using the **UDDI References > List** option.
  - *serviceKey* is the service key that the UDDI registry has assigned to the service. Here is an example of a UDDI service key: 34280367-0ECF-46CE-B804-14C21D6D0FB1. For more information about UDDI service keys, see Publishing a Web service to a UDDI registry for deployment to the gateway.
- Target Service Name.** If the Web service WSDL contains more than one service, or the WSDL is located through a UDDI registry, type the target service name from the target service WSDL.
  - Target Service Namespace.** If the Web service WSDL contains more than one service, or the WSDL is located through a UDDI registry, type the namespace of the target service name from the target service WSDL.
  - Target Service Identity Information.** Type the identity by which the target service is known within the Web services gateway. This identity need not be unique.

**Note:** If you later add more target services to this gateway service, and also write a filter to select a target service, you might use the Target Service Identity Information to select a particular target service from the set.

- If you want to publish this gateway service to one or more UDDI registries (selected in a previous step), enter the UDDI business key in the field provided in the **UDDI Publication Properties** section. This key identifies the business category under which you want your service to appear in UDDI. To get a list of valid business keys, look up businesses in a UDDI registry. Here is an example of a UDDI business key: 08A536DC-3482-4E18-BFEC-2E2A23630526. For more information about UDDI business keys see Publishing a Web service to a UDDI registry for deployment to the gateway.
- Click **OK**.

If the processing completes successfully, the list of deployed Web services is updated to include the new Web service. Otherwise, an error message is displayed.

After deployment, use the list deployed Web services option to add multiple target services for this gateway service, or to change the following resources that the Web service uses:

- Channels.
- Filters
- **5.1 +** JAX-RPC handlers
- UDDI references
- Security bindings.

If you enabled the “Authorization Policy - Control access to this service” check box , you must now enable Web service operation-level authorization.

**Data type representation - Choosing between Generic classes and Deployed Java classes:** When you deploy a Web service, you use the **Message part representation** option to choose between **Generic classes** and **Deployed Java classes**.

As your message passes through the gateway, the message parts are represented as actual Java objects. These are the data type definitions that are used for each part:

- One of the set of XML schema and SOAP supported simple and compound data types, or
- A complex type defined in the WSDL schema section.

A complex type is a data type represented by a Java class (such as a user-written class) that is not part of the native Java language.

**Generic classes** and **Deployed Java classes** can both represent simple, compound and complex data types, subject to the following constraints:

- The gateway only supports the simple and compound types that are listed in Web services gateway - Supported types.
- Only **Deployed Java classes** can represent vectors, enumerations, hashtables and maps (but either **Generic classes** or **Deployed Java classes** can represent arrays).
- If **Deployed Java classes** are used to represent complex types, then the actual Java classes representing these complex types must be deployed to the application server on which the gateway is running.
- If the target service uses Java or EJB WSDL bindings (that is, if the target service is a Java class deployed on the local application server, or it is an enterprise bean) then you must use **Deployed Java classes**, and you must make the bindings available as described in Deploying Web services with Java bindings.

The gateway schema parser determines all top-level types that are defined in the WSDL schema section, and generates mappings to **generic classes** for all of these types. This mapping enables the gateway to forward requests (and responses) containing most complex data type parameters (and return values) to a remote destination without requiring the actual Java classes representing these complex types to be deployed to the application server on which the gateway is running. If your Web service uses complex data types, and there is no other constraint that forces you to use **Deployed Java classes**, then select **Generic classes**.

Performance is the same whether you choose to use **Deployed Java classes** or **Generic classes**.

**Complex data types - Mapping namespaces to packages:** If you write your own WSDL file that describes your Web service (rather than use an automated tool such as WebSphere Studio Application Developer) and the service uses complex data types, then follow these guidelines for mapping namespaces to packages.

For working with complex data types, there is no industry-wide standard way of mapping namespaces to packages (in fact the Java API for XML-based Remote Procedure Call (JAX-RPC) standard states that the tools must make up their own standard). This is the standard used by the gateway:

- Set the namespace of the complex data type to the Java package name of the class.
- Set the complex data type name to the name of the class.

For example, if the Java class you are using for the complex data type is `random.RandomData`, then the namespace of the complex data type is `random`, and the complex data type name is `RandomData`. And if the full package name is `com.ibm.www.random`, then the namespace is `www.ibm.com/random`.

See also the troubleshooting tip about working with Web services that use complex data types.

### **Deploying Web services with Java bindings:**

For Web services deployed with Java bindings (or EJB bindings even where the Web service is on a different server) you must make additional classes available to the gateway.

For **EJB bindings**, make the EJB client JAR file available. If the Web service is deployed on the same server as the gateway, the necessary interfaces and classes are already visible. If not, implement one of the following options:

- Copy the EJB client JAR file into the `install_root/lib` or `install_root/lib/app` directory, where `install_root` is the root directory for your installation of IBM WebSphere Application Server (by default `WebSphere/AppServer`).
- Update the application server class path to include the EJB client JAR file.

For **Java bindings**, make the Java classes for the Web service available by implementing one of the following options:

- Copy the JAR file or class files that contain the Java classes into the *install\_root/lib* or *install\_root/lib/app* directory.
- Update the application server class path to include the JAR file.
- Wrap the Java classes in an enterprise bean and deploy it on the same application server. WebSphere Application Server then makes the classes available to the gateway application.

**Web services gateway - Supported types:** When you deploy a Web service, you use the **Message part representation** option to choose between Generic classes and Deployed Java classes. This topic gives reference information about the data types that are supported in each case.

## Simple types

The following table gives a list of the XML schema (and Java equivalent) simple types that are supported by both **Generic classes** and **Deployed Java classes**:

XML schema (and Java equivalent) simple type
xsd:string
xsd:float
xsd:double
xsd:int
xsd:boolean
xsd:byte
xsd:short
xsd:long
xsd:decimal
xsd:QName
xsd:date
xsd:timeInstant

## Compound types

From a Java perspective, compound types are types with constituent elements. These elements are either identified purely by name (for example, a Java class with several member properties) or by ordinal position (for example, a List data structure like Array or Vector).

The following table gives a list of the SOAP (and Java equivalent) compound types that are supported:

SOAP compound type	Java equivalent compound type	Supported by
Array	Java array	Generic classes and Deployed Java classes
Vector	java.util.Vector	Deployed Java classes
Vector	java.util.Enumeration	Deployed Java classes
Map	java.util.Hashtable	Deployed Java classes
Map	java.util.Map	Deployed Java classes

## Complex types

A complex type is a data type represented by a Java class (such as a user-written class) that is not part of the native Java language.

Complex types can include combinations of simple types, compound types, and other complex types. For example, the children of a complex type might be represented by another complex type, or by any of the simple or compound types.

**Generic classes** and **Deployed Java classes** can both represent complex types, subject to the constraints described in Data type representation - Choosing between Generic classes and Deployed Java classes.

### ***Deploying a gateway service for proxy operation:***

Use this task to set the gateway to act purely as a proxy for your service, and use JAX-RPC handlers to set the endpoints for incoming request messages for the service.

For an overview of proxy operation, see JAX-RPC handlers and proxy operation. To deploy a gateway service for proxy operation, complete the following steps:

1. Create a JAX-RPC handler that sets the `transport.url` context property to set the target endpoint for the service. It is up to you to decide how you want to establish the target endpoint, and to create and configure the `transport.url` JAX-RPC handler. For information about creating JAX-RPC handlers, see the IBM developerWorks article *A developer introduction to JAX-RPC, Part 2: Mine the JAX-RPC specification to improve Web service interoperability*.
2. Deploy the new handler to the gateway.
3. Deploy the SOAP over HTTP channel (if not already deployed). Proxy operation is only supported by this channel.
4. Deploy your Web service. Ensure that you set the following parameters:
  - a. Enable **Selective SOAP parsing**.
  - b. Enable **Operate in Proxy Mode**.
  - c. Optional: If you want to override the default proxy service template WSDL, then type the location of your replacement template WSDL file in the **Proxy Service Template WSDL Override** field.
  - d. Select the SOAP over HTTP channel from the list of available channels.

**Note:** Do *not* set any **Target Service Properties**.

5. List the Web services that are deployed to the Web services gateway, then complete the following steps to add the handler that sets the target endpoint for this service:
  - a. Click the name of your service in the list. A form is displayed that you can use to view and modify the current deployment details.
  - b. Scroll down to **Channels**, then click **Edit JAX-RPC handler configuration**.
  - c. Select the handler that sets the target endpoint for this service, and add it to the handler chain.

Ensure that all service requesters append an `operationMode` HTTP query string parameter to the Web address for the gateway service. The value of this parameter is either `oneway` or `requestResponse`. For example, if the Web address of the deployed proxy service is `http://.../wsgwssoap/urn:wsgw::ProxyService`, then requesting clients indicate that they are sending a one-way request by using the following URL:

```
http://.../wsgwssoap/urn:wsgw::ProxyService?operationMode=oneway
```

### ***Publishing a Web service to a UDDI registry for deployment to the gateway:***

The gateway interacts with UDDI registries as described in UDDI registries - Web service directories that integrate with the Web services gateway. When you deploy a Web service to the gateway, you enable UDDI interaction by entering a UDDI reference, and (depending upon what you are trying to do) either or both of the following pieces of information:

- The service key that the UDDI registry has assigned to this service.
- The business key that identifies the UDDI business category under which you want your service to appear in the UDDI registries.

You get these two keys from the UDDI registry. To help you understand what UDDI service keys and business keys are, and where you find them in a UDDI registry, this topic describes how to publish a Web service to a UDDI registry.

The following description of the UDDI publication process requires that you specify a Technical model. Technical models are generic categories. Using these models, a UDDI registry user can search for a type of service, rather than needing to know the access details for a specific service. The gateway makes no use of Technical model information because it only interacts with UDDI registries at the level of specific Web services.

The following task steps include specific navigation instructions. These instructions describe how you publish a Web service to the IBM WebSphere UDDI Registry. If you are working with a different UDDI registry, then the specific navigation is different but the underlying principles are the same.

1. Specify a business:
  - a. To get a list of valid business keys, look up businesses in the UDDI registry. Here is an example of a UDDI business key: 08A536DC-3482-4E18-BFEC-2E2A23630526.
  - b. If you do not find an appropriate existing business in the UDDI registry, then use the **Add a business** option on the **Advanced Publish** section of the Publish pane to add a new one.
2. Add a technical model:
  - a. Select **Add a technical model** on the **Advanced Publish** section of the Publish pane.
  - b. Enter the name as specified for the target namespace of your binding (or interface) WSDL file, then add a description (if required).
  - c. Add a category of Type unspsc and value wsdlSpec (the Key name field can be left blank).
  - d. Add an overview URL specifying the Web address for your binding WSDL file, then add a description (if required).

**Note:** The binding and the service definition for your Web service might be held in separate WSDL files, therefore be careful to type the Web address of the WSDL file that defines the *binding*.

  - e. Click **Publish Technical Model**.
3. Add a service:
  - a. Select **Show owned entities** on the **Advanced Publish** section of the Publish pane.
  - b. Select **Add a Service** for your business.
  - c. Enter the name as specified for the target service in your WSDL file, then add a description (if required).
  - d. For the **Access point** verify that the correct Web address type is selected (for example http for an HTTP access point), then enter the value of the soap:address location (or its equivalent) from your service definition WSDL file (for example http://yourhost:80/SimpleTest/servlet/rpcrouter).
  - e. For the **Technical model** select **Add**, then find the required Technical model by entering a suitable prefix and selecting **Find technical models**, then enable the check box for the required Technical model and click **Update**.
  - f. Click **Publish Service**.

The UDDI registry assigns a service key to your service, and publishes the service.



## Removing Web services from the Web services gateway

To remove a Web service, complete the following steps:

1. Display the Web services gateway administrative user interface.
2. In the navigation pane, click the following link:

### Services

- [Remove](#)

The main pane is updated with a list of all the Web services that are deployed to the Web services gateway. Next to each entry in the list is a check box.

3. Select the check box for every Web service that you want to remove.
4. Click **OK**.

If the processing completes successfully, the list of deployed Web services is updated. Otherwise, an error message is displayed.

If the Web service that you want to remove is also published to one or more UDDI registries, then the internal WSDL file is required to remove the service from the gateway. Therefore if the service is published to UDDI, and the processing does not complete successfully, check that the WSDL file is still available at the location defined for the service in the **Target Services -> WSDL Location**. For more information, see Listing and managing gateway-deployed Web services.

---

## Running the Web services gateway Samples

The following pre-built Samples are available for use with the Web services gateway:

- The standard Stock Quote service Sample, that requires an Internet connection.
- The Address Book service Sample, that allows the storing and retrieval of names and addresses.

These Samples, and documentation on how to use them, are available through the Web services gateway Samples link on the Samples Central page of the DeveloperWorks WebSphere Web site.

If you want to test the gateway taking service definitions from a private UDDI registry such as the IBM WebSphere UDDI Registry, complete the following additional steps:

1. Publish the WSDL for each of these Samples to UDDI. (For more information, see Publishing a Web service to a UDDI registry for deployment to the gateway and the documentation for your private UDDI registry).
2. Instruct the gateway to locate the service through the UDDI registry, as described in Deploying Web services to the Web services gateway.

---

## Passing SOAP messages with attachments through the Web services gateway

The Web services gateway supports Web services that pass attachments in a MIME message.

This support is included in the SOAP over HTTP channel.

Attachments are carried through the various gateway components and passed to the target service. The content MIME type of each attachment is preserved.

When the target service is deployed to a Java API for XML-based Remote Procedure Call (JAX-RPC) compliant server, you can access the attachments on the target service using the `javax.activation.DataHandler` handler.



The WSDL representing a SOAP messages with attachments service must define the attachment parts in the binding section. The `mime:multipartRelated`, `mime:part` and `mime:content` tags are used to describe the attachment.

For more information, see the following topics:

- SOAP messages with attachments - a definition.
- Writing the WSDL extensions for SOAP messages with attachments.

The following scenarios are not supported:

- Using the Apache SOAP channel.
- Using DIME.
- Using the `mime:mimeXml` WSDL tag.
- Nesting a `mime:multipartRelated` inside a `mime:part`.
- Using arrays or vectors of `DataHandlers`, images, and so on.

The MIME headers from the incoming message are not preserved for referenced attachments. The outgoing message contains new MIME headers for `Content-Type`, `Content-Id` and `Content-Transfer-Encoding` that are created by WSIF.

## SOAP messages with attachments - a definition

From an architecture and external specification viewpoint, SOAP Messages with Attachments is an extension to the SOAP 1.1 Recommendation from the World Wide Web Consortium (W3C).

The W3C SOAP messages with attachments document describes a standard way to associate a SOAP message with one or more attachments in their native format (for example GIF or JPEG) by using a multipart MIME structure for transport. This document defines specific use of the Multipart/Related MIME media type, and rules for the use of URI references to refer to entities bundled within the MIME package. It outlines a technique for a SOAP 1.1 message to be carried within a MIME multipart/related message in such a way that the SOAP processing rules for a standard SOAP message are not changed.

An associated W3C document Web Services Description Language (WSDL) outlines a technique for including bindings to MIME types in a WSDL file.

## Writing the WSDL extensions for SOAP messages with attachments

The following example WSDL illustrates a simple operation that has one attachment called `attch`:

```
<binding name="MyBinding" type="tns:abc" >
 <soap:binding style="rpc" transport="http://schemas.xmlsoap.org/soap/http"/>
 <operation name="MyOperation">
 <soap:operation soapAction=""/>
 <input>
 <mime:multipartRelated>
 <mime:part>
 <soap:body parts="part1 part2 ..." use="encoded" namespace="http://mynamespace"
 encodingStyle="http://schemas.xmlsoap.org/soap/encoding"/>
 </mime:part>
 <mime:part>
 <mime:content part="attch" type="text/html"/>
 </mime:part>
 </mime:multipartRelated>
 </input>
 </operation>
</binding>
```

In this type of WSDL extension:

- There must be a `part` attribute (in this example `attch`) on the input message for the operation (in this example `MyOperation`). There can be other input parts to `MyOperation` that are not attachments.

- In the binding input there must either be a `<soap:body>` tag or a `<mime:multipartRelated>` tag, but not both.
- For MIME messages, the `<soap:body>` tag is inside a `<mime:part>` tag. There must only be one `<mime:part>` tag that contains a `<soap:body>` tag in the binding input and that must not contain a `<mime:content>` tag as well, because a content type of `text/xml` is assumed for the `<soap:body>` tag.
- There can be multiple attachments in a MIME message, each described by a `<mime:part>` tag.
- Each `<mime:part>` tag that does not contain a `<soap:body>` tag contains a `<mime:content>` tag that describes the attachment itself. The type attribute inside the `<mime:content>` tag is not checked or used by the gateway. It is there to suggest to the application using the gateway what the attachment contains. Multiple `<mime:content>` tags inside a single `<mime:part>` tag means that the backend service expects a single attachment with a type specified by one of the `<mime:content>` tags inside that `<mime:part>` tag.
- The `parts="..."` attribute inside the `<soap:body>` tag is assumed to contain the names of all the SOAP parts in the message, but not the attachment parts. If there are only attachment parts, then you must specify `parts=""` (empty string). If you omit the `parts` attribute altogether, then the gateway assumes ALL parts including the attachments - which means the attachments will appear twice.

In your WSDL you might have defined a schema for the attachment (for instance as a `binary[]`). The gateway silently ignores this mapping and treats the attachment as a Data Handler.

You do not need to mention unreferenced attachments in the WSDL bindings.

---

## Developing Web services gateway extensions

Use this task to develop your own extensions for the gateway.

**Note:** **5.1+** The gateway system level interfaces described in this task and its sub-tasks are deprecated. However, you can achieve similar results using Java API for XML-based Remote Procedure Call (JAX-RPC) handlers. For more information, see [JAX-RPC handlers - An alternative to gateway filters](#).

The filters that you create as described in this topic and sub-topics use the gateway `WorkArea` to maintain context. The gateway `WorkArea` is only available with IBM WebSphere Application Server Enterprise Version 5 or WebSphere Business Integration Server Foundation Version 5.1. Therefore you can only deploy these filters on WebSphere Application Server Enterprise Version 5 or WebSphere Business Integration Server Foundation Version 5.1.

To extend the functionality of the gateway, you can write your own Java programs. The gateway does not provide any application programming interfaces, but there are system level interfaces that you can use. Specific guidance on using these interfaces is provided in the following topics:

- [Writing a filter for the Web services gateway.](#)
- [Using a filter to select a target service and port.](#)
- [Capturing Web service invocation information from the Web services gateway.](#)
- [Handling exceptions for the Web services gateway.](#)

For additional technical details of the Web services gateway, see the Javadoc.

## Writing a filter for the Web services gateway

**Note:** **5.1+** Use of the `Filter` interface is deprecated. However, you can achieve similar results using Java API for XML-based Remote Procedure Call (JAX-RPC) handlers. For more information, see [JAX-RPC handlers - An alternative to gateway filters](#).

The filters that you create as described in this topic use the gateway `WorkArea` to maintain context. Specifically, they extend `FilterImpl` which implements this for you:

```

public int getContextVersion() {
 return Filter.CONTEXT_VERSION_WORKAREA;
}

```

The gateway WorkArea is only available on WebSphere Application Server Enterprise Version 5 or WebSphere Business Integration Server Foundation Version 5.1.

To use this information you should be familiar with using a Java 2 platform, Enterprise Edition (J2EE) session bean development environment such as IBM WebSphere Studio Application Developer.

A Web services gateway filter is essentially a J2EE session bean implementing specific Home and Remote interfaces.

To write a filter using IBM WebSphere Studio Application Developer, for deployment on WebSphere Application Server Enterprise Version 5 or WebSphere Business Integration Foundation Version 5.1, complete the following steps. For more detailed information on writing session beans, see the WebSphere Studio documentation topic *Developing enterprise beans - overview*.

1. Open the J2EE perspective.
2. To create a new EJB application project, complete the following steps:
  - a. Select **File > New > Enterprise Application Project**. The Project Creation wizard opens.
  - b. In the Project Creation wizard, complete the following steps:
    - 1) Select the version of the J2EE specification that you want to use, then click **Next**.

**Note:** **5.1+** If you want to enable the local client view (so that the gateway can make local EJB calls to your filter, as described in *Tuning the Web services gateway*) then you must specify J2EE version 1.3 or later.

- 2) Type your project name.
- 3) Create a new module project for the EJB project only. Clear the other **Application Client Project**, **Web Project** and **Connector Project** check boxes as necessary.
- 4) Clear the **Web module** check box.
- 5) Click **Finish**.

Your new EJB application project and associated EAR project are created.

3. To add the extra JAR files that your EJB module needs that are not already in the enterprise application server `/lib` directory, complete the following steps:
  - a. Select **File > Import**.
  - b. In **Select an import source**, select **File system** then click **Next**.
  - c. In the Import window, complete the following steps:
    - 1) Select `wsgw_root/client` as the source directory.  
where `wsgw_root` is the root directory for your installation of the gateway.
    - 2) Select the `wsgwejb.jar` file.
    - 3) Select the root directory of your new EAR project as the **destination for imported resources**.
    - 4) Click **Finish**.

**Note:** You might see one or more instances of the following WebSphere Studio warning message. You should ignore these messages.

```

IWAE0024W The Manifest Class-Path for archive wsgwejb.jar contains
an entry, name.jar, that is not resolveable to a file or module in
the EAR earName.

```

- d. Repeat the previous **File > Import** process to add any other extra JAR files that your EJB module needs.
- e. In the J2EE Hierarchy view, from the pop-up menu for your EJB module, select **Open With > JAR Dependency Editor**.
- f. In the JAR Dependencies window, select all the JAR files listed.

- g. Close the JAR Dependencies window, then click **Yes** in the Save Resource window to save your changes.
- 4. To add extra JAR files to the Java build path for your EJB module, complete the following steps:
  - a. In the J2EE Hierarchy view, select your EJB module **Properties**.
  - b. In the Properties window, verify that the following JAR files are included on the **Java Build Path**:
    - *install\_root/lib/jrom.jar*
    - *install\_root/lib/qname.jar*
    - *install\_root/lib/wsd14j.jar*
    - *install\_root/lib/wsif.jar*
    - *wsgw\_root/client/wsgwejb.jar*

where *install\_root* is the root directory for your installation of IBM WebSphere Application Server Enterprise Version 5 or IBM WebSphere Business Integration Foundation Version 5.1, and *wsgw\_root* is the root directory for your installation of the gateway.

- c. Add any other JAR files or projects that you need for compiling your filter.
- d. Click **OK**.
- 5. To create the session bean, complete the following steps:
  - a. Select **File > New > Enterprise Bean**. The Enterprise Bean Creation wizard opens.
  - b. In the Enterprise Bean Creation wizard, complete the following steps:
    - 1) Select your EJB project, then click **Next**.
    - 2) Verify that **Session Bean** is selected.
    - 3) Enter a name for the bean.
    - 4) Enter a suitable package name for the bean.
    - 5) Click **Next**.
  - c. In the Enterprise Bean Details window, complete the following steps:
    - 1) Accept the defaults offered for **Session type** (Stateless) and **Transaction type** (Container).
    - 2) Accept the defaults offered for **Bean supertype** (<none>), **Bean class** and **EJB binding name**.
    - 3) **5.1 +** If you selected J2EE version 1.3 or later for your enterprise application project, and you want the gateway to make local EJB calls to your filter, then enable the **Local client view** by completing the next two steps.
    - 4) **5.1 +** For the **Local client view: Local Home Interface**, click **Class...** then select the `com.ibm.wsgw.beans.FilterLocalHome` interface.
    - 5) **5.1 +** For the **Local client view: Local Interface**, click **Class...** then select the `com.ibm.wsgw.beans.FilterLocal` interface.

**Note:** When the FilterLocal interface is used, the arguments to the filter methods are passed by reference. You should ensure that the `filterRequest()` and `filterResponse()` methods do not modify their arguments, unless the modification is intended.

    - 6) For the **Remote client view: Remote Home Interface**, click **Class...** then select the `com.ibm.wsgw.beans.FilterHome` interface.
    - 7) For the **Remote client view: Remote Interface**, click **Class...** then select the `com.ibm.wsgw.beans.FilterRemote` interface.
    - 8) Click **Next**.
  - d. In the EJB Java Class window, specify the **Bean superclass** as `com.ibm.wsgw.beans.FilterImpl`, then click **Finish**.

Your new session bean is created.

- 6. The generated Java code for your session bean does not implement the filter. To update the code, complete the following steps:
  - a. In the J2EE Hierarchy view, expand your session bean to show Java code entries for the Home interface, the Remote interface and for the session bean.

- b. In the J2EE Hierarchy view, double-click the entry for the session bean code. In the editor view, the generated code opens for editing.
- c. In the editor view, add the following import statements:

```
import com.ibm.wsgw.*;
import com.ibm.wsgw.beans.*;
import org.apache.wsif.*;
import java.rmi.RemoteException;
```

If you selected J2EE version 1.3 or later for your enterprise application project, then omit the `import java.rmi.RemoteException` statement from the set of import statements added here.

- d. Select **File > Save** to save the file. Ignore any errors.

If you selected J2EE version 1.3 or later for your enterprise application project, then your filter methods must not throw a `java.rmi.RemoteException` exception. For J2EE version 1.3 or later:

- Omit the statement `import java.rmi.RemoteException` from the set of import statements added in the previous step.
- Remove all instances of `throws RemoteException` from your filter methods.
- Code your filter to override the `getContextVersion()` method from `FilterImpl` (to eliminate the `RemoteException` from this method). Specifically, implement a `getContextVersion()` method in your filter to return `Filter.CONTEXT_VERSION_WORKAREA` as shown in the following example:

```
public int getContextVersion()
{
 return Filter.CONTEXT_VERSION_WORKAREA;
}
```

7. To add the unimplemented methods of the `Filter` interface to your session bean, complete the following steps:
  - a. Open the Outline view (select **Window > Show View > Outline**).
  - b. In the Outline view, from the pop-up menu for your session bean, select **Override Methods**.
  - c. In the Override Methods window, select all the `Filter` methods to override then click **OK**.

The methods of the `Filter` interface are added to your session bean.

8. Select **File > Save** to save the file. Any errors from the previous **File > Save** are resolved.
9. Develop your filter.

The exact steps that you take to develop your filter depend upon what you want it to do. However, to develop any filter, you use the following resources:

- The `Filter` interface.
- The gateway Javadoc for the `Filter` interface.
- The additional information on the `Filter` interface that is in *Web services gateway - The Filter interface*
- The gateway message context. (This contains the context values for each message that comes into the gateway. Your filter acts upon these values.)
- The gateway Javadoc for the `GatewayContextNames` class. (To use the gateway message context values, you import the `GatewayContextNames` class.)
- The additional information on the gateway message context values that is in *Web services gateway - The message context values*.
- The gateway `WorkArea`. (Filters use this resource to get and set the message context values, as described in *Web services gateway - The Filter interface*.)
- The WSIF Javadoc for the following WSIF objects:
  - `WSIFRequest`
  - `WSIFResponse`
  - `WSIFMessage`
  - `WSIFException`

(The methods of the `Filter` interface use these objects, as described in *Web services gateway - The Filter interface*.)

- The following example code.

**Note:** You must observe the J2EE programming model, and ensure that any non-gateway services you use are available on all platforms on which the filter might be expected to run. For example, do not use static variables to store state information because on certain platforms, or in certain configurations (such as a scaled gateway deployment), a filter might be invoked in a different Java Virtual Machine (JVM) for each request.

This example shows you how to access the context and get values in the filterRequest method of a filter.

```
import javax.naming.InitialContext;
import javax.naming.NamingException;

import com.ibm.websphere.workarea.UserWorkArea;
import com.ibm.websphere.workarea.WorkAreaException;

import com.ibm.wsgw.GatewayContextNames;

...

try
{
 // Lookup the WorkArea gateway context in JNDI
 InitialContext ctx = new InitialContext();
 UserWorkArea wsgwContext =
 (UserWorkArea)ctx.lookup("services:websphere/WSGW/workarea");

 // Get the currently selected port name
 String Ptype =
 (wsgwContext.get(GatewayContextNames.TARGET_PORT_NAME)).getClass().getName();
 String ThePortname =
 (String) wsgwContext.get(GatewayContextNames.TARGET_PORT_NAME);

 // Get the currently selected target service WSDL location
 String Xtype =
 (wsgwContext.get(GatewayContextNames.TARGET_SERVICE_LOCATION)).getClass().getName();
 TargetServiceLocation WSDLObject =
 (TargetServiceLocation) wsgwContext.get(GatewayContextNames.TARGET_SERVICE_LOCATION);

 String ServiceLocation = WSDLObject.serviceLocation;
 int ServiceLocationType = WSDLObject.serviceLocationType;
 String ServiceName = WSDLObject.serviceName;
 String ServiceNamespace = WSDLObject.serviceNamespace;
}
catch (NamingException e)
{
 // Handle any exceptions thrown by the InitialContext here
}
catch (WorkAreaException e)
{
 // Handle exceptions thrown by UserWorkArea here
}

...
```

After you have developed your filter, you need to generate deployment code and export the enterprise application. To do this using IBM WebSphere Studio Application Developer, complete the following steps:

1. Open the J2EE perspective.
2. In the J2EE Hierarchy view, from the pop-up menu for your EJB module, select **Generate > Deploy and RMIC code**.
3. In the Generate Deploy and RMIC Code window, select the beans for which you want to generate code, then click **Finish**.
4. To configure the deployment descriptor properties for your bean, complete the following steps:
  - a. In the J2EE Hierarchy view, from the pop-up menu for your bean, select **Open With > EJB Deployment Descriptor**.



- b. On the **Beans** tab, set the Java Naming and Directory Interface (JNDI) name to the Filter class name. This name is used as the Home Location when the filter is deployed to the gateway.
  - c. Close the EJB Deployment Descriptor window, then click **Save** to save the changes.
5. To export the enterprise application, complete the following steps:
    - a. In the J2EE Hierarchy view, from the pop-up menu for your project, select **Export ...**.
    - b. In the **Select an export destination:** box, select **EAR file**, then click **Next**.
    - c. In the **What resources do you want to export?** box, type your project name.
    - d. In the **Where do you want to export resources to?** box, type the destination directory.
    - e. Select any other options that you require, then click **Finish**.

You are now ready to install your filter into WebSphere Application Server Enterprise Version 5 or WebSphere Business Integration Server Foundation Version 5.1 (as described in the next to last step of Installing the gateway into a deployment manager cell and Installing the gateway into a stand-alone application server), then deploy your filter.

## Web services gateway - The Filter interface

**Note:** **5.1 +** Use of the Filter interface is deprecated. However, you can achieve similar results using Java API for XML-based Remote Procedure Call (JAX-RPC) handlers. For more information, see [JAX-RPC handlers - An alternative to gateway filters](#).

This topic gives more information on using each of the methods of the Filter interface. It supplements the information provided in the following Javadoc:

- The gateway Javadoc for the Filter interface.
- The WSIF Javadoc for
  - WSIFRequest
  - WSIFResponse
  - WSIFMessage
  - WSIFException

The Filter interface represents an object that is called during service invocation. You can register a bean that implements this interface to be called just before a request invocation, or just after a response is received for a particular service.

Gateway filters use the gateway WorkArea to get and set the gateway message context values. You get the gateway WorkArea partition from `services:websphere/WSGW/workarea` in the Java Naming and Directory Interface (JNDI), then use the UserWorkArea API to get and set data within the message context.

These are the two main methods you use for developing your filter:

- `FilterAction filterRequest(org.apache.wsif.WSIFRequest request, org.apache.wsif.WSIFResponse response)`.
- `FilterAction filterResponse(org.apache.wsif.WSIFRequest request, org.apache.wsif.WSIFResponse response)`.

Another important method that requires a specific value to be returned is `getContextVersion()`.

If you want your filter to change the WSIFResponse and WSIFRequest messages, then note that changes to messages are only recognized if the `setUpdatedRequest` and `setUpdatedResponse` methods are called on the returned FilterAction object. The FilterAction object can also dictate whether or not to continue processing the message by calling the `setContinueProcess` method.



## **filterRequest method**

The filterRequest method is called by the gateway Manager just before a request is sent to a target service. The return value from the method can indicate not to send the request.

The request parameter contains the request WSIFMessage. This consists of a set of named parts. Each part has a value that is encoded as an instance of an appropriate Java object. Filters can change the values of the Java object instances, but should not add or remove parts, or replace the values of parts with ones of a different type.

The Filter interface might decide not to proceed with the request. In that case the interface has three options:

- Throw a FilterException exception. The gateway logs the exception but continues processing filters and the request invocation.
- Throw a WSGWException exception. The gateway logs and rethrows the exception, and processing of filters and the request is stopped. The exception then goes back to the receiving channel, and the channel must determine what to do with the exception (in the case of SOAP-based channels, this results in a Fault message back to the client). This should only be done for unexpected errors in the filter.
- Return a FilterAction object with the continueProcessing flag set to false. In this case the response message in the FilterAction object can also be set, and is sent to the originator of the request. No further filters are invoked.

If the request or response is modified, then it must be returned in an instance of the FilterAction class. If this return is not done, any change to the response is ignored by the gateway.

## **filterResponse method**

The filterResponse method is called by the gateway Manager just after a response is received from a target service. The response parameter contains the response or fault WSIFMessage. This message consists of a set of named parts.

Each part has a value that is encoded as an instance of an appropriate Java object. Filters can change the values of the Java object instances, but should not add or remove parts, or replace the values of parts with ones of a different type.

The Filter interface might decide not to proceed with the response. In that case the interface has three options:

- Throw a FilterException exception. The gateway logs the exception but continues processing filters and the response invocation.
- Throw a WSGWException exception. The gateway logs and rethrows the exception, and processing of filters and the response is stopped. The exception then goes back to the receiving channel, and the channel must determine what to do with the exception (in the case of SOAP-based channels, this results in a Fault message back to the client). This should only be done for unexpected errors in the filter.
- Return a FilterAction object with the continueProcessing flag set to false. In this case the fault WSIF Message is set in the response message in the FilterAction object, and is sent to the originator of the request. No further filters are invoked.

If the response is modified, then it must be returned in an instance of the FilterAction class. If this return is not done, any change to the response is ignored by the gateway.

If the filter throws a FilterException exception, it is logged but the gateway continues to process other filters. If the filter throws a WSGWException exception, processing of the response is stopped.

See also Creating and returning a SOAP fault message from a filterResponse method.

## getVersionString method

The getVersionString method returns a string form of the filter implementation version. The gateway uses this form when logging events relating to the filter so that the exact version of the filter implementation is known.

## getContextVersion method

The getContextVersion method indicates the approach that this filter uses to access context information. To access the message context information for WebSphere Application Server Enterprise Version 5 or WebSphere Business Integration Server Foundation Version 5.1, implement this method to return the value: Filter.CONTEXT\_VERSION\_WORKAREA.

## init method

The init method tells the filter that it is configured with the gateway. This method is called when you deploy a filter to the gateway.

## destroy method

The destroy method tells the filter that it is no longer configured with the gateway. This method is called when you remove a filter from the gateway.

## Creating and returning a SOAP fault message from a filterResponse method

This topic provides an example of how to create and return a SOAP fault message from a GatewayFilter.filterResponse method.

**Note:** **5.1+** Use of the Filter interface is deprecated. However, you can achieve similar results using Java API for XML-based Remote Procedure Call (JAX-RPC) handlers. For more information, see [JAX-RPC handlers - An alternative to gateway filters](#).

```
public FilterAction filterResponse(WSIFRequest wsifRequest, WSIFResponse wsifResponse)
 throws FilterException, WSGWException, RemoteException {
 // Construct the fault message
 WSIFMessage faultMessage = new WSIFDefaultMessage();
 faultMessage.setObjectPart(WSIFConstants.SOAP_FAULT_ACTOR, "mySoapFaultActor");
 faultMessage.setObjectPart(WSIFConstants.SOAP_FAULT_CODE, "mySoapFaultCode");
 faultMessage.setObjectPart(WSIFConstants.SOAP_FAULT_STRING, "mySoapFaultString");
 faultMessage.setObjectPart("stackTrace", "myStackTraceDetails");
 faultMessage.setObjectPart("otherDetails", "myOtherDetails");
 // repeat faultMessage.setObjectPart("aaaa", "bbbb");
 // for each additional detail element
 // Set the fault message into the wsifResponse object
 wsifResponse.setFaultMessage(faultMessage);
 wsifResponse.setIsFault(true);

 // Return the updated response in the filterAction object
 FilterAction filterAction = new FilterAction();
 filterAction.setUpdatedResponse(wsifResponse);
 filterAction.setContinueProcessing(false);
 return filterAction;
}
```

The previous example works with both the Apache SOAP and the SOAP over HTTP channels. If you use the SOAP over HTTP channel, the previous example returns the following SOAP fault response:

```
<?xml version="1.0" encoding="UTF-8"?>
<soapenv:Envelope xmlns:soapenv="http://schemas.xmlsoap.org/soap/envelope/"
 xmlns:xsd="http://www.w3.org/2001/XMLSchema"
 xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance">
 <soapenv:Body>
```

```

<soapenv:Fault>
 <faultcode xmlns:ns1="http://xml.apache.org/axis/">ns1:mySoapFaultCode</faultcode>
 <faultstring>mySoapFaultString</faultstring>
 <faultactor>mySoapFaultActor</faultactor>
 <detail>
 <stackTrace>myStackTraceDetails</stackTrace>
 <otherDetails>myOtherDetails</otherDetails>
 </detail>
</soapenv:Fault>
</soapenv:Body>
</soapenv:Envelope>

```

## Web services gateway - The gateway message context values

The gateway message context holds the contextual information for each incoming message.

In WebSphere Application Server Enterprise Version 5 or WebSphere Business Integration Server Foundation Version 5.1 you can use the **Context Field Constant** values if you import the `com.ibm.wsgw.GatewayContextNames` class.

For basic information on the fields that are available in the context, see the Javadoc for the `GatewayContextNames` class. Additional information on all of these fields except `AUTH_SUBJECT` and copyright is provided in the following table.

Treat all of the context fields as Read only. If your filter attempts to write to a context field, you do not get an error message (because the write does not actually fail) but subsequent system behavior is not readily predictable.

If you want to change the target service location and port name fields, then you can use the Routing interface to get the list of valid target service locations and to select the target service location. For more information see Using a filter to select a target service and port.

Context Field Constant	Description
<code>AUTH_PASSWORD</code>	Read the password from the incoming HTTP request (where available)
<code>AUTH_USER_NAME</code>	Read the user name from the incoming HTTP request (where available).
<code>GATEWAY_SERVICE_NAME</code>	Name of the gateway service for which the request was received.
<code>MESSAGE_ID</code>	A server-unique ID that is set by the channel and used to correlate messages, for example in trace. You can make this ID globally unique by using the host name as a prefix.
<code>RECEIVING_CHANNEL_NAME</code>	Name of the channel on which the request was received.
<code>REQUEST_SOAP_HEADERS</code>	Retrieve the SOAP headers for an inbound SOAP request. SOAP Headers are returned as a vector of nodes.
<code>RESPONSE_SOAP_HEADERS</code>	Retrieve the SOAP headers for an outbound SOAP response. SOAP Headers are returned as an array of nodes.
<code>RETRY_COUNT</code>	Number of retries for the request. This constant is NOT CURRENTLY USED.
<code>SOAP_HEADERS</code>	Retrieve the SOAP headers for an inbound SOAP request. SOAP Headers are returned as a vector of nodes. This constant is deprecated (use the <code>REQUEST_SOAP_HEADERS</code> constant).

Context Field Constant	Description
TARGET_PORT_NAME	Currently selected port name. This constant is set after service invocation, therefore the constant can only be got by response filters. See also Using a filter to select a target service and port.
TARGET_SERVICE_LOCATION	Gives the location of the currently selected target service WSDL, name and namespace.
TIMEOUT_TIME	Time-out value for the response. This constant is NOT CURRENTLY USED.

## Using a filter to select a target service and port

Use this task to write a filter to select a target service and port.

**Note:** **5.1+** Use of the Routing interface, and of filters, is deprecated. However, you can achieve similar results using Java API for XML-based Remote Procedure Call (JAX-RPC) handlers. For more information, see *JAX-RPC handlers - An alternative to gateway filters*.

When a request is received by the gateway, it must determine what the target service really is, and what port to use to access that service.

The gateway represents each target service as a gateway service. Each gateway service can map to one or more target services, but without filters there is no point in mapping multiple targets because the gateway always picks the first one. If you want to map multiple targets, you also need to write pluggable filters that are configured for each gateway service and that can select the target service.

You write a filter as described in *Writing a filter for the Web services gateway*. Your filter can get the list of potential target services from the Routing service. It needs to select the target service, then call the Routing service to set the target service. This selection clears any prior selection of a target service port. Your filter might also use the Routing service to select the target port for the service invocation.

The Routing service provides a non-standard interface which is defined in the topic *The **Routing** interface*. The implementation of the Routing service is not pluggable.

The Home object for this service must implement the `com.ibm.wsgw.beans.RoutingHome` interface and be located in the Java Naming and Directory Interface (JNDI) in the `websphere/WSGW/Routing` directory.

This is the sequence of events for a filter to determine and set the target service:

1. The filter is called with a `WSIFRequest` request.
2. The filter obtains the list of potential target services from the Routing service.
3. The filter selects the target service.
4. The filter calls the Routing service to set the target service.

Each target service is identified by the unique location of the target service definition, and by the target service identity information (which might not be unique). To select the target service, your filter can either get the table of mappings from target service location to identity information, then choose a target service to use; or it can call the `setSelectedTargetServiceIdentity` method with the required identity string (relying on the definition of the target service identity information). The routing service then selects the first target service it finds for the current gateway service that matches the specified identity information (using `String.equals()`).

When you use routing to set the target service or the target port, the Routing service updates the request context. Because the request context has changed, you then need to set the request object in the `FilterAction` object that you return from the `filterRequest` method (see the filter interface).

## Web services gateway - The Routing interface

**Note:** **5.1+** Use of the Routing interface is deprecated. However, you can achieve similar results using Java API for XML-based Remote Procedure Call (JAX-RPC) handlers. For more information, see JAX-RPC handlers - An alternative to gateway filters.

This topic gives more information on using each of the methods of the Routing interface. It supplements the information given for this interface in the gateway Javadoc.

The Routing interface encapsulates a service that manages routing for requests. Filters can use this service to select the target service and port.

The set methods all return a WSIFRequest object that contains the updated context information.

To get information on the currently selected target service, use the following Routing interface methods:

### **getSelectedTargetServiceLocation**

This method returns the currently selected target service location for the request.

### **getTargetPortName**

This method returns the currently selected target port name for the request.

**Note:** In this version of the gateway, this method always returns blank.

### **getTargetServiceDefinition**

This method returns the currently selected target service definition for the request.

To set the target port, use the following Routing interface method:

### **setTargetPortName**

This method sets the selected target port name for the request.

To get information on all potential target services, use the following Routing interface method:

### **getTargetServices**

This method returns the set of target service names that are mapped by the gateway service on which the request was received.

To set the target service, use either of the following Routing interface methods:

### **setSelectedTargetServiceLocation**

This method sets the selected target service location for the request. The selected port name is reset by this call.

### **setSelectedTargetServiceIdentity**

This method sets the selected target service identity for the request. The target service identity does not need to be unique, therefore the first target service found with matching identity information is set. If no service is found that matches, the method throws a WSGWException exception.

## Capturing Web service invocation information from the Web services gateway

**Note:** **5.1+** Use of the MessageWarehouse interface is deprecated. However, you can achieve similar results using Java API for XML-based Remote Procedure Call (JAX-RPC) handlers. For more information, see JAX-RPC handlers - An alternative to gateway filters.

The Web services gateway has not implemented a service that stores operational messages, but the gateway does contain an interface, the MessageWarehouse interface, to encapsulate such a service. This interface is driven by channels on receipt of requests.

If you have your own system for handling (classifying, storing and retrieving) operational messages, you can use this system to log the gateway operational messages through the MessageWarehouse interface.

The Home object for this service must implement the `com.ibm.wsgw.beans.MessageWarehouseHome` interface and be located in the Java Naming and Directory Interface (JNDI) in the `websphere/WSGW/MessageWarehouse` directory.

## Web services gateway - The MessageWarehouse interface

**Note:** **5.1+** Use of the MessageWarehouse interface is deprecated. However, you can achieve similar results using Java API for XML-based Remote Procedure Call (JAX-RPC) handlers. For more information, see [JAX-RPC handlers - An alternative to gateway filters](#).

This topic provides more information on each method of the MessageWarehouse interface. It supplements the information given for this interface in the gateway Javadoc.

The MessageWarehouse interface encapsulates a service which stores messages for archiving. Channels use this interface to log incoming requests for the purposes of non-repudiation.

A default implementation of this interface is not provided by the Web services gateway. If no implementation is present, the interface is not used.

### logRequest

This method stores a request, with information about the channel and the originator of the request.

The channel calls this method when a request is received, after the user is authenticated and the message decrypted. The channel might provide information to identify the originator of the request, and to identify the channel. The request is logged as a `WSIFMessage` message.

You can also log additional information regarding receipt of the request (for example any associated digital certificates) as `Serializable` objects.

### logResponse

This method stores a response, with information about the channel and the destination of the response.

The channel calls this method when a response is about to be sent, before the response is encrypted. The channel might provide information to identify the destination of the response, and to identify the channel. The response is logged as a `WSIFMessage` message.

You can also log additional information regarding sending of the response (for example any associated digital certificates) as `Serializable` objects.

### logException

This method stores a request, with information about the channel and the originator of the request, when an exception is thrown to the channel during request processing. This method allows the exact request and exception information to be logged before the channel decides the actions to take.

You can also log additional information regarding request and exception (for example any associated digital certificates) as `Serializable` objects.

## Handling exceptions for the Web services gateway

Use this task to help you capture information on exception handling activities for the gateway.

**Note:** **5.1+** Use of the `ExceptionHandler` interface is deprecated.

During normal processing of a Web service invocation, the target service might generate a fault message and pass it to the channel for forwarding to the originator. As far as the gateway is concerned there is no difference between processing a normal output message and processing a fault message.

When an exception occurs during processing of a request, the channel needs some way to decide what to do with the exception. A service is needed that provides a pluggable handler that can look at the message, the exception and other information to decide whether to throw the exception back to the originator, or to construct a fault message.

This service is not provided with the Web services gateway, but the gateway does contain an interface to encapsulate such a service. The **ExceptionHandler** interface allows channels to call an exception handling service, and report the exceptions to a third party for analysis.

The Home object for this service must implement the `com.ibm.wsgw.beans.ExceptionHandlerHome` interface and be located through the Java Naming and Directory Interface (JNDI) in the `websphere/WSGW/ExceptionHandler` directory.

## Web services gateway - The ExceptionHandler interface

**Note:** **5.1+** Use of the ExceptionHandler interface is deprecated.

This topic provides more information on using each of the methods of the ExceptionHandler interface. This documentation supplements the information provided in the gateway Javadoc for the ExceptionHandler interface.

The ExceptionHandler interface encapsulates a service that channels use to react to exceptions that occur during request and response processing in the gateway.

### **handleException**

This method is called by a channel when an exception is caught as a result of message processing.

The return value indicates the action for the channel to take. The actions include:

- Re-throw the original exception.
- Throw a new exception (The handler throws this exception).
- Convert the exception into a fault message.

If no exception handler is installed, the MessageWarehouse (if any) logs the exception, then the exception is rethrown.

---

## Scaling the Web services gateway

You cannot install the gateway into an application server that is part of a cluster. However, you can still scale the gateway by installing an identical gateway on several application servers, then using a plug-in to the front-end HTTP server to balance the load.

The servers that you include in the load-balancing group can be stand-alone application servers, or application servers that are members of a deployment manager cell, or a mixture of the two. Each server in the group can be hosted on a separate node, or all the servers can be hosted on one or two nodes.

If you create each application server in the group as a member of a single deployment manager cell (that is, create all the servers on machines that are nodes of a single cell), you can use the deployment manager tooling for the cell to help you install and update the gateway installation on each application server in the group, no matter where each server is physically hosted.



You configure and maintain an instance of the gateway that is separate from your deployment group as the source configuration for every gateway in the group. You use the back up option to save the source configuration as a **Shared** configuration, then you restore the saved configuration to every gateway in the group.

**5.1** To scale the Web services gateway, complete the steps given in *Scaling the Web services gateway* using unlocked configuration deployment.

## Scaling the Web services gateway using unlocked configuration deployment

You scale the gateway by installing an identical gateway on several application servers to create a scaling group, then using a plug-in to the front-end HTTP server to balance the load across the group.

**5.1** If any (or all) of the application servers in your scaling group are hosted on nodes that are members of a deployment manager cell, the cell cannot also contain other gateway servers that are not part of the scaling group.

You configure an instance of the gateway that is separate from the scaling group, then maintain this gateway as the source configuration for every gateway in the group. You use the back up and restore options to save the source configuration as a **Shared** configuration, then restore the saved configuration to every gateway in the group.

To scale the gateway using unlocked configuration deployment, complete the following steps:

1. Select or create each member of your scaling group of application servers. If you install WebSphere Application Server, include the gateway files in the installed configuration either by accepting the “typical install” option, or by choosing the “custom install” option **Web services > Web services gateway**.
2. Optional: Add the nodes for the application servers in the scaling group to a deployment manager cell. (See “Managing nodes” in the Information Center.) Note that the cell cannot also contain other gateway servers that are not part of the scaling group.
3. Install a gateway on every application server in the scaling group.
4. Select and configure the gateway that is the source configuration for every gateway in the scaling group. Ensure that your source configuration meets the following requirements:
  - This gateway is not hosted on an application server that is part of the group.
  - The WSDL URI for exported definitions is set to indicate the HTTP server that is used by the scaling group.
  - The HTTP channels deployed to this gateway have their end point addresses set to indicate the HTTP server that is used by the scaling group.
  - **5.1** The HTTP channels deployed to this gateway have their end point addresses configured as two-part addresses (see *Deploying channels to the Web services gateway*).
5. Use the back up and restore options to save the source configuration as a **Shared** configuration, then restore the saved configuration to every gateway in the scaling group.
6. Use the WebSphere Application Server caching proxy and load balancer in the front-end HTTP server to balance the load by sharing incoming HTTP requests among the application servers in the scaling group. For detailed information see the information center for WebSphere Application Server Edge components. This information center contains complete documentation for the caching proxy and the load balancer in the following PDF online books:
  - *WebSphere Application Server Concepts, Planning, and Installation for Edge Components*.
  - *WebSphere Application Server Caching Proxy Administration Guide*.
  - *WebSphere Application Server Programming Guide for Edge Components*.

7. Apply all subsequent gateway updates only to the source configuration gateway. This is particularly important with regard to Web services that are published to UDDI by the gateway. When you save a gateway configuration, the record of which Web services are gateway-published to UDDI is not saved (for reasons that are explained in Backing up and restoring UDDI publication links). If you change the gateway that you use as the source for configuration updates, you lose this record.

---

## Securing the Web services gateway

The gateway provides facilities for secure communication between the service requester and the gateway, and between the gateway and the target service.

To configure these facilities for each deployed Web service, see [Enabling Web Services Security \(WS-Security\)](#) for the gateway.

The gateway provides basic authentication and authorization facilities based upon the broader security features of WebSphere Application Server. See [Enabling basic authentication and authorization](#) for the gateway.

The gateway can also invoke Web services that include `https://` in their addresses, if you have configured your Java and WebSphere Application Server security properties accordingly. To check your security property settings, see [Invoking Web services over HTTPS](#).

In many installations, the gateway requires access to the Internet for invoking external Web services. To configure the gateway to work with authenticating proxies, see [Enabling proxy authentication](#) for the gateway.

For hints on solving security-related problems, see [Troubleshooting the Web services gateway](#).

## Enabling Web Services Security (WS-Security) for the gateway

You can configure the gateway for secure transmission of SOAP messages using tokens, keys, signatures and encryption in accordance with the Web Services Security (WS-Security) draft recommendation.

For more information on how WS-Security is implemented in WebSphere Application Server Network Deployment, see [Securing Web services](#). For more information on the approach taken by the gateway to implementing this emerging standard, see [The Web services gateway and WS-Security](#).

The gateway sits between the service requester (the client) and the target Web service. You configure the gateway to act as the target service from the point of view of the client, and as the client from the point of view of the target service. You therefore need to get, from the owning parties, the WS-Security configurations for both the client and the Web service. This information is found in the following files on the owners systems:

- Key stores (.ks and .jceks files).
- Certificate stores (.cer files).
- Security settings (the `ibm-webservicesclient-ext.xml` file for the client, and the `ibm-webservices-ext.xml` file for the Web service).
- Binding information - for example the location of a keystore file on the file system (the `ibm-webservicesclient-bnd.xml` file for the client, and the `ibm-webservices-bnd.xml` file for the Web service).

If the client is hosted on WebSphere Application Server, and the Web service security settings are created using IBM Web services tooling (for example WebSphere Studio Application Developer), then the files that contain the security settings and binding information have the exact file names (\*.xml) noted previously. For clients and Web services from other vendors, these files have different file names.

You need to copy the key store and certificate store files to the gateway file system, and to enter and configure for the gateway the security settings that are contained in the .xmi files. The security settings are entered and configured manually using the gateway administrative user interface. There are tools available (for example WebSphere Studio Application Developer) that can parse the .xmi files for you.

You use the **Gateway > Security** option to configure the security bindings (the tokens, keys, signatures and encryption methods) that are available to the gateway, as described in Configuring the gateway security bindings. You then configure the level of security that applies at each stage of the transmission (and note that different levels of security, including no security, can be applied at each stage):

- From the service requester to the gateway.
- From the gateway to the target service.
- From the target service back to the gateway.
- From the gateway back to the service requester.


For information on how to configure the security levels, see the following topics:

- Editing the service security configuration - how to configure secure communication for this gateway service between the service requester (the client) and the gateway.
- Editing the target service security configuration - how to configure secure communication between the gateway and the target service.

### The Web services gateway and WS-Security

You can configure the gateway for secure transmission of SOAP messages using tokens, keys, signatures and encryption in accordance with the emerging Web Services Security (WS-Security) specification.

In a normal (non gateway) WS-Security scenario, the message flows are as shown in the following figure:

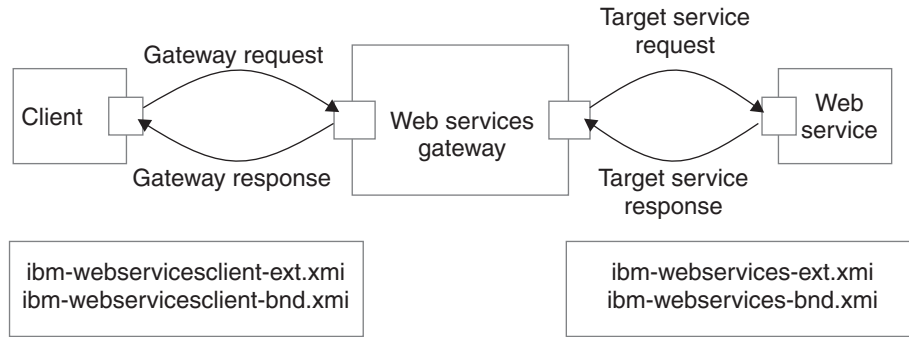


The client generates a request that is handled by the client Web services engine. This engine reads the client security configuration and applies the security defined in the `ibm-webservicesclient-ext.xmi` file to the SOAP message. It gets additional binding information from the `ibm-webservicesclient-bnd.xmi` file (for instance, the location of a keystore on the file system).

On receipt of a SOAP message, the Web services engine on the server refers to the \*.xmi files for the called Web service. In this case, the `ibm-webservices-ext.xmi` file tells the engine what security the incoming message must have (for example, that the body of the message must be signed). If the message does not comply, then it is rejected. The Web services engine verifies any security information, then passes the message on to the Web service that is called.

On the response from server to client, the process is reversed. The Web service \*.xmi files tell the Web services engine what security to apply to the response message, and the client \*.xmi files tell the client engine what security to require in the response message.

When the gateway is introduced, the scenario is more complex. You can think of this scenario as two separate request and response invocations, client to gateway and gateway to target service, as shown in the following figure:



In this scenario, the client application and the Web service are unchanged, and still have the same security settings in their \*.xmi files. However, the gateway is unsecured. Secure SOAP messages cannot travel through the gateway unchanged, and must be processed on receipt. The gateway therefore needs to act as the target service from the point of view of the client, and as the client from the point of view of the target service. In this scenario the security settings for the Web service need to be configured for the view of the service that the gateway presents to the client, and the security settings for the associated gateway target services (remember that there might be multiple target services deployed for a single gateway service) need to be configured with the security settings for the client.

WS-Security settings for the gateway are configured manually using the gateway administrative user interface.

## Configuring the gateway security bindings

You use the **Gateway > Security** option to configure all the security bindings (the tokens, keys, signatures and encryption methods) that are available to the gateway. This information describes the security bindings that are used to secure the SOAP messages that pass between service requesters (clients) and the gateway, and between target services and the gateway.

You receive this security binding information direct from the service requester or target service provider, in the form of an `ibm-webservicesclient-bnd.xmi` file for the client, and an `ibm-webservices-bnd.xmi` file for the Web service. You extract the information from these .xmi files, then manually enter it into the gateway security bindings forms that are described in the following steps.

1. Display the Web services gateway administrative user interface.
2. In the navigation pane, click the following link:

### Gateway

- [Security](#)

The **Configure Gateway Security Bindings** form is displayed.

This form is divided into eight sections, one for each type of security binding. The following comments apply to every section:

- To add a new binding, click **Add new binding type**.
- To amend an existing binding, click on the name of the binding.
- To delete an existing binding, click **remove** next to the name of the binding.
- If you choose to add or amend a binding, then a binding information subform is displayed. Within this form:
  - Help is provided in comments on the subform, and in hover-help next to each field.
  - Fields marked on the subform with an asterisk (\*) are required fields.
  - For some bindings, the fields that are required are different depending on whether you are *adding* or *amending* a binding.
  - Many fields are populated by making a selection from a list.
  - Information entered in one binding information subform often appears in a list in another subform.

- When you have finished entering information for a binding, click **OK**.

3. Add, amend or delete **Signing Information**.

The *signing information* specifies the configuration for digital signature validation and message signing.

If you choose to add or amend a signing information binding, then the **Define Signing Information** form is displayed. In this form you enter a name for the binding (if you are adding a new binding), and the following binding details:

- Signature method
- Digest method
- Canonicalization method
- Signing key name
- Signing Key Locator
- Certificate Path (a choice between trusting *any* certificates, and trusting certificates from a particular *Trust Anchor* and *Certificate Store*).

**Note:** The signing information can also be used for X.509 certificate validation when **Authentication Method** is IDAssertion and **ID Type** is X509Certificate in the gateway service configuration. In such cases, you must fill in the Certificate Path fields only.

4. Add, amend or delete **Encryption Information**.

The *encryption information* specifies the configuration for encrypting and decrypting messages.

If you choose to add or amend an encryption information binding, then the **Define Encryption Information** form is displayed. In this form you enter a name for the binding (if you are adding a new binding), and the following binding details:

- Key locator
- Encryption key name
- Key encryption algorithm
- Data encryption algorithm

5. Add, amend or delete **Trust Anchors**.

A *trust anchor* specifies a list of key store configurations that contain root trusted certificates.

If you choose to add or amend a trust anchor, then the **Define Trust Anchor** form is displayed. In this form you enter a name for the trust anchor (if you are adding a new anchor), and the following details:

- Key store type
- Key store path
- Key store password

**Note:** These configurations are used for certificate path validation of the incoming X.509-formatted security tokens. You create the keystore using the Java command-line tool `keytool`. ( See "Managing digital certificates" in the Information Center.) The `ikeymantool` is not supported.

6. Add, amend or delete **Certificate Stores**.

A *certificate store* specifies a list of untrusted, intermediate certificate files. It is used for certificate path validation of incoming X.509-formatted security tokens.

If you choose to add or amend a certificate store, then the **Define Certificate Store** form is displayed. In this form you enter a name for the certificate store (if you are adding a new store), and the following details:

- Certificate Store Provider

When you amend an existing certificate store, you are given an extra option to add or remove X.509 certificates from the list of certificates that are contained within this store. When you add an X.509 certificate, you specify the full path for the certificate.

7. Add, amend or delete **Key Locators**.

A *key locator* specifies a configuration that is used to retrieve keys for signature and encryption. A key locator class can be customized to retrieve keys from other types of repositories. The default implementation retrieves keys from a keystore.

If you choose to add or amend a key locator, then the **Define Key Locators** form is displayed. In this form you enter a name for the key locator (if you are adding a new key locator), and the following details:

- Classname
- Key store type
- Key store path
- Key store password

When you amend an existing key locator, you are given two extra options:

- Add or remove key entries. For each additional key, you specify:
  - Key name
  - Key alias
  - Key password

**Note:** You do not need to list all the certificate entries as keys; instead, the distinguished name (DN) of the certificate is used as the search key.

- Add or remove additional properties for the configuration. For each additional property, you specify:
  - Property name
  - Property value

8. Add, amend or delete **Trusted ID Evaluators**.

A *trusted ID evaluator* determines whether the identity (ID)-asserting authority is trusted.

If you choose to add or amend a trusted ID evaluator, then the **Define Trusted ID Evaluator** form is displayed. In this form you enter a name for the trusted ID evaluator (if you are adding a new evaluator), and the following details:

- Class name

When you amend an existing trusted ID evaluator, you are given an extra option to add or remove additional properties for the configuration. For each additional property, you specify:

- Property name
- Property value

9. Add, amend or delete **Login Mappings**.

A *login mapping* specifies a configuration for validating security tokens within incoming messages.

If you choose to add or amend a login mapping, then the **Define Login Mapping** form is displayed. In this form you enter a name for the login mapping (if you are adding a new mapping), and the following details:

- Authentication method
- JAAS Configuration name
- Callback handler factory class name
- Token type URI
- Token type local name
- **5.1 +** Nonce maximum age
- **5.1 +** Nonce clock skew

When you amend an existing login mapping, you are given extra options to add or remove additional properties for the *configuration*, and additional properties for the *callback handler factory*. For each additional property, you specify:

- Property name
- Property value

10. Add, amend or delete **Login Bindings**.

A *login binding* specifies a configuration for generating security tokens within outgoing messages.



If you choose to add or amend a login binding, then the **Define Login Binding** form is displayed. In this form you enter a name for the login binding (if you are adding a new binding), and the following details:

- Authentication method
- Callback handler class name
- Token type URI
- Token type local name
- Basic authentication user ID
- Basic authentication password

When you amend an existing login binding, you are given extra options to add or remove additional properties for the *configuration*, and additional properties for the *callback handler*. For each additional property, you specify:

- Property name
- Property value

## Editing the service security configuration

Before you can select the security settings that are applied for an individual Web service, you must configure the gateway security bindings.

For each Web service, you can select the security settings that are applied between the service requester (the client) and the gateway. These settings are specified for each stage of the transmission:

- From the service requester to the gateway (the client request).
- From the gateway back to the service requester (the gateway response).

You receive this security settings information from the service requester and from the target service provider in the following form:

- An `ibm-webservicesclient-bnd.xmi` file for the client, and an `ibm-webservices-bnd.xmi` file for the Web service, from which you process the security bindings information as described in *Configuring the gateway security bindings*.
- An `ibm-webservicesclient-ext.xmi` file for the client, and an `ibm-webservices-ext.xmi` file for the Web service, which contain the information on the levels of security (integrity, confidentiality and identification) that are required when this Web service exchanges messages with a service requester. These levels are also the settings that the gateway needs to apply when it makes the equivalent gateway service available to a service requester.

To set the security settings that are applied between the service requester (the client) and the gateway, complete the following steps:

1. List the gateway-deployed Web services
2. Click the name of a Web service in the list. A form is displayed that you can use to view and modify the current deployment details for this Web service, and add or remove multiple target services.
3. In the **Service Security** section, select the **Edit service security configuration** option.

The service security configuration form is displayed. This form is divided into the following sections:

- **Gateway Security Properties** (the Actor URI)
- **Client Request Security Properties** (integrity, confidentiality and identification settings)
- **Gateway Response Security Properties** (the response Actor URI, and integrity and confidentiality settings)
- **Security bindings** (request bindings and response bindings).

The following comments apply to every section:

- Help is provided in comments on the form, and in hover-help next to each field.
- There are no required fields.
- Many fields are populated by making a selection from a list.

4. In the **Gateway Security Properties** section, set the Actor URI.

If you specify an Actor URI, then only SOAP security headers with this Actor URI are processed.



5. In the **Client Request Security Properties** section, set the following security levels:
  - a. Set the **Integrity** level. Set the parts of the incoming SOAP message that must be signed (the Body, the Timestamp and the Security Token).
  - b. Set the **Confidentiality** level. Set the parts of the incoming SOAP message that must be encrypted (the Body and the Username Token).
  - c. Set the **Identification** level. Set the identification methods that will be accepted (Basic Authentication , Digital Signature, ID Assertion and Lightweight Third Party Authentication (LTPA)). See "Configuring Lightweight Third Party Authentication" in the Information Center.
6. In the **Gateway Response Security Properties** section, set the following security levels:
  - a. Set the Response Actor URI.
 

**Note:** If you specify a Response Actor URI, then the SOAP security header in the response message has this Actor URI.
  - b. Set the **Integrity** level. Set the parts of the response SOAP message that must be signed (the Body and the Timestamp).
  - c. Set the **Confidentiality** level. Set whether or not the Body of the response SOAP message must be encrypted.
7. In the **Security bindings** section, set the Request bindings and the Response bindings to use. You choose these bindings (Signing Information, Encryption Information, Trusted ID Evaluator and Login Mappings) from pull-down lists. The available items in these lists are those that you previously entered, as described in Configuring the gateway security bindings.
8. When you have finished editing the service security configuration, click **Apply changes**.

## Editing the target service security configuration

Before you can select the security settings that are applied for a target Web service, you must configure the gateway security bindings.

For each target Web service, you can select the security settings that are applied between the target Web service and the gateway. These settings are specified for each stage of the transmission:

- From the gateway to the target service (the target service request).
- From the target service back to the gateway (the target service response).

You receive this security settings information from the service requester and from the target service provider - usually in the following form:

- An `ibm-webservicesclient-bnd.xmi` file for the client, and an `ibm-webservices-bnd.xmi` file for the Web service, from which you process the security bindings information as described in Configuring the gateway security bindings.
- An `ibm-webservicesclient-ext.xmi` file for the client, and an `ibm-webservices-ext.xmi` file for the Web service, which contains the information on the levels of security (integrity, confidentiality and identification) that are required when this Web service exchanges messages with a service requester. These settings are also needed by the gateway when it calls the target service on behalf of the service requester.

To set the security settings that are applied between the target service and the gateway, complete the following steps:

1. List the gateway-deployed Web services
2. Click the name of a Web service in the list. A form is displayed that you can use to view and modify the current deployment details for this Web service, and add or remove multiple target services.
3. In the **Target Services** section, click the name of a target Web service in the list. A form is displayed, containing the same fields that you filled in when you added the target service, and also the following additional fields:
  - **Started** (a check box).

- **Enable target service security** (a check box).
  - **Edit target service security configuration.**
4. Select the **Edit target service security configuration** option.  
The target service security configuration form is displayed. This form is divided into the following sections:
    - **Target Service Security Properties** (the Actor URI)
    - **Target Service Request Security Properties** (the Target Actor URI, and the integrity, confidentiality and identification settings)
    - **Target Service Response Security Properties** (integrity and confidentiality settings)
    - **Security bindings** (request bindings and response bindings).
 The following comments apply to every section:
    - Help is provided in comments on the form, and in hover-help next to each field.
    - There are no required fields.
    - Many fields are populated by making a selection from a list.
  5. In the **Target Service Security Properties** section, set the Actor URI.  
If you specify an Actor URI, then only SOAP security headers with this Actor URI are processed.
  6. In the **Target Service Request Security Properties** section, set the following security levels:
    - a. Set the Target Actor URI.  
If you specify a Target Actor URI, then the SOAP security header in the request message has this Actor URI.
    - b. Set the **Integrity** level. Set the parts of the outgoing SOAP message that must be signed (the Body, the Timestamp and the Security Token).
    - c. Set the **Confidentiality** level. Set the parts of the outgoing SOAP message that must be encrypted (the Body and the Username Token).
    - d. Set the **Identification** level. Set the identification methods that will be accepted (None, Basic Authentication, Digital Signature, ID Assertion and Lightweight Third Party Authentication (LTPA)). See "Configuring Lightweight Third Party Authentication" in the Information Center.
  7. In the **Target Service Response Security Properties** section, set the following security levels:
    - a. Set the **Integrity** level. Set the parts of the response SOAP message that must be signed (the Body and the Timestamp).
    - b. Set the **Confidentiality** level. Set whether or not the Body of the response SOAP message must be encrypted.
  8. In the **Security bindings** section, set the Request bindings and the Response bindings to use.  
You choose these bindings (Signing Information, Encryption Information and Login Binding) from pull-down lists. The available items in these lists are those that you previously entered, as described in Configuring the gateway security bindings.
  9. When you finish editing the target service security configuration, click **Apply changes**.

## Enabling basic authentication and authorization for the gateway

In addition to the security options described in Enabling Web Services Security (WS-Security) through the gateway, you can also use the broader security features of WebSphere Application Server to enable basic authentication and authorization.

Basic authentication can be applied at two levels, as described in the following topics:

1. Enabling gateway-level authentication.
2. Enabling Web service operation-level authorization.

For **gateway-level authentication**, you set up a role and realm for the gateway on WebSphere Application Server Web server and servlet container, and define the user ID and password that is used by the gateway to access the role and realm. You also modify the gateway channel applications so that they

only give access to the gateway to service requesters that supply the correct user ID and password for that role and realm. This means that gateway-level authentication must be enabled before you install any channels.

For **operation-level authorization**, you apply security to individual methods in a Web service. To do this, you create an enterprise bean with methods matching the Web service operations. These EJB methods perform no operation and are just entities for applying security. Existing WebSphere Application Server authentication mechanisms can be applied to the enterprise bean. Before any Web service operation is invoked, a call is made to the EJB method. If authorization is granted, the Web service is invoked. Your target Web service is protected by wrapping it in an EAR file, and applying role-based authorization to the EAR file. This process is explained in general terms in *Operation-level security - role-based authorization*.

**Note:**

- If you want to enable operation-level authorization, you must first enable gateway-level authentication.
- After gateway-level authentication is enabled, filters have access to the authentication information for the requester.

The Web services gateway can also invoke Web services that include `https://` in their addresses, if the Java and WebSphere security properties have been configured to allow it. To check your security property settings, see *Invoking Web services over HTTPS*.

For hints on solving security-related problems, see *Troubleshooting the Web services gateway*.

## Enabling gateway-level authentication

A number of default gateway-level authentication settings are included in the gateway. There is a default role of `AuthenticatedUsers` which includes the special group `AllAuthenticatedUsers`. When security is enabled, you must supply a user ID and password to use the gateway administrative interface or invoke a gateway service.

This task covers the three main areas in which you might want to make changes:

- Changing the default gateway-level authentication settings.

If you want to change the default gateway-level authentication settings, you must do so before you install any channels. When you run the script that installs the gateway (*either* into an application server that is part of a deployment manager cell *or* into a stand-alone application server) you also install the following channels:

- Apache SOAP channel 1.
- SOAP over HTTP channel 1.

If you change the default gateway-level authentication settings after you install the gateway, you then need to run the gateway installation program again.

- Enabling gateway-level authentication.

You can enable gateway-level authentication at any time. After gateway-level authentication is enabled, filters have access to the authentication information for the requester.

- Assigning users and groups to roles.

You can assign users and groups to roles at any time.

To enable gateway-level authentication, complete the following steps:

1. To change the default gateway-level authentication settings, choose a tool that is available for your version of WebSphere Application Server:

- **5.1+** The WebSphere Application Server Assembly Toolkit.

Use this tool to complete the following steps:

- a. Set up a role and realm for the gateway on the WebSphere Application Server Web server and servlet container.
  - b. Define the user ID and password that are used by the gateway to access the role and realm.
  - c. Modify the gateway channel applications so that they only give gateway access to service requesters that supply the correct user ID and password for that role and realm.
2. To enable gateway-level authentication, complete the following steps:
    - a. Start the WebSphere Application Server administrative server.
    - b. Start the administrative console.
    - c. In the navigation pane, select **Security > Global Security**.
    - d. In the main pane, on the **Configuration** tab, select the Enabled check box.
    - e. Save the settings.
    - f. Stop then restart the application server.
    - g. Close the administrative console.
  3. To map users to roles using the administrative console, complete the following steps:
    - a. Start the WebSphere Application Server administrative server.
    - b. Start the administrative console.
    - c. In the navigation pane, select **Application > Enterprise Applications > wsgw**. In the main pane, an option to map security roles to users and groups appears in the Additional Properties table.
    - d. Modify the security roles and save the settings.
    - e. Repeat the previous two steps for each enterprise application that you want to modify.
    - f. Stop then restart the application server.
    - g. Close the administrative console.

For more information, see "Assigning users and groups to roles" in the Information Center.

**Note:** The current Jacl installation scripts do not let you assign users to roles as part of installing the gateway into an application server that is part of a deployment manager cell or into a stand-alone application server.

You can now enable operation-level authorization, or install the gateway.

## Enabling operation-level authorization

Use this task to apply security to individual methods in a Web service.

Before you begin this task you must first enable gateway-level authentication.

You can only apply operation-level authorization to a Web service that is already deployed to the gateway with the check box "Authorization Policy - Control access to this service" enabled.

For operation-level authorization you create an enterprise bean with methods matching the Web service operations. These EJB methods perform no operation and are just entities for applying security. You can apply existing WebSphere Application Server authentication mechanisms to the enterprise bean. Before any Web service operation is invoked, a call is made to the EJB method. If authorization is granted, the Web service is invoked.

Your target Web service is protected by wrapping it in an EAR file (*your\_webservice.ear*), then applying role-based authorization to the EAR file. This process is explained in general terms in Operation-level security - role-based authorization. The *your\_webservice.ear* file is then imported into the *wsgwauth.ear* file and the *wsgwauth.ear* file is modified to set the roles and assign them to methods. The modified *wsgwauth.ear* file is then deployed in WebSphere Application Server, and users are assigned to the previously-defined roles.

The wsgwauth.ear file contains an EAR file for each Web service that you protect. The installation version of the wsgwauth.ear file is in *install\_root/installableApps*, where *install\_root* is the root directory for your installation of IBM WebSphere Application Server (by default WebSphere/AppServer). For the first Web service that you protect through operation-level authorization, you copy the installation version of the wsgwauth.ear file and store your copy outside of the application server file system. For each subsequent Web services that you protect, you further modify the same copy of the wsgwauth.ear file.

To enable Web service operation-level authorization, complete the following steps for each Web service that you want to protect:

1. For the first Web service that you protect, make your own copy of the *install\_root/installableApps/wsgwauth.ear* file in a convenient location outside of the application server file system.
2. To create the *your\_webservice.ear* file, complete the following steps:
  - a. Open a command prompt.
  - b. Go to the *install\_root/WSGW/scripts/auth* directory.
  - c. Enter the following command:

```
WSGWAAuthGen location your_webservice
```

where:

- *location* is the Web address for the gateway. This must include the root context.
- *your\_webservice* is the name of the service as deployed in the gateway. This is case-sensitive.

For example

```
WSGWAAuthGen http://host:port/wsgw AddressBook
```

where *host* and *port* are the host name and port number for the application server on which the gateway is installed.

The Web service name and operation name can contain characters (such as a dash (-), period (.) and ampersand (&)) that are disallowed in an EJB class name and method name. Therefore these are translated during the generation process of the *your\_webservice.ear* file. A message appears informing you of any name changes.

The *your\_webservice.ear* file is created in the current directory. There is also a temporary directory *current\_directory/ejb* that you can delete.

3. To finish assigning roles and protecting methods, complete the steps given in the following topics:
  - **5.1+** Using the Assembly Toolkit to enable operation-level authorization.
4. To install the modified copy of the wsgwauth.ear file, complete the following steps:
  - a. Check that the modified wsgwauth.ear file is stored in the convenient location outside of the application server file system that you chose in step 1. Keep the wsgwauth.ear file in this location for subsequent reuse and further modification.
  - b. Start the WebSphere Application Server administrative console.
  - c. In the navigation pane, select **Applications > Install an Application**.
  - d. Use **Install New Application** to install the modified copy of the wsgwauth.ear file. Select the users or groups to assign to the roles when prompted.

**Operation-level security - role-based authorization:** When you build an EAR file, you can define roles and apply them to methods. When you deploy the EAR file, you can assign individual users or groups to roles. You can use this feature of EAR files to add role-based security to your Web service. For example:

1. You have a Web service that controls access to important information, and you want to give read-only access to some users, and write access to others.
2. When you build the EAR file you define two roles: READ and WRITE. You apply the READ role to the getData method and the WRITE role to the writeData method.

3. When you deploy the EAR file in WebSphere Application Server, you assign All Authenticated Users to the READ role and individual users to the WRITE role.
4. When a user tries to access `WebService.getData`, their user name and password is checked by the operating system or by Lightweight Third Party Authentication (LTPA). See "Configuring Lightweight Third Party Authentication" in the Information Center.

**Using the Assembly Toolkit to enable operation-level authorization:**

This task assumes that you have already completed the initial steps for Enabling operation-level authorization.

As is explained in general terms in Operation-level security - role-based authorization, your target Web service is protected by wrapping it in an EAR file and applying role-based authorization to the EAR file. In this task, the EAR file that contains your Web service (*your\_webservice.ear*) is imported into *thewsgwauth.ear* file (which contains all of the protected Web services) and the *wsgwauth.ear* file is modified to set the roles and assign them to methods. This modified *wsgwauth.ear* file is then deployed in WebSphere Application Server and users are assigned to the previously defined roles.

Use the Assembly Toolkit component of the Application Server Toolkit to complete the following steps:

1. Start the Assembly Toolkit then open the J2EE perspective.
2. From the File menu select **File > Import > EAR**, then browse to select your copy of the *wsgwauth.ear* file. Two J2EE projects are created:
  - *wsgwauth*
  - *wsgwauthejb*
3. From the File menu select **File > Import > EAR**, then browse to select the *your\_webservice.ear* file. Two J2EE projects are created:
  - *your\_webservice*
  - *your\_webserviceejb*
4. Select the *your\_webserviceejb* project, then edit the **EJB Deployment Descriptor**. For every security role that you want to create, repeat the following steps:
  - a. On the Assembly Descriptor tab, add the required security role (for example **READER**).
  - b. Follow the on-screen instructions to assign one or more method permissions to the security role.
  - c. Save your changes.
5. To import the *your\_webservice.ear* file into the *wsgwauth.ear* file, complete the following steps:
  - a. Select the *wsgwauth* project, then edit the **EAR Deployment Descriptor**.
  - b. On the Module tab, add the *your\_webserviceejb* enterprise bean from the *your\_webserviceejb* project.
  - c. Save your changes.
6. To ensure that the authorization enterprise bean can reference the newly-imported enterprise bean, complete the following steps to add an EJB reference:
  - a. Select the *wsgwauthejb* project, then edit the **EJB Deployment Descriptor**.
  - b. On the References tab, select **Add EJB Reference**.
  - c. In the Add EJB Reference pane, click **Browse**.
  - d. Select the **Enterprise bean in a different EJB project** radio button, then browse to select the *your\_webserviceejb* enterprise bean.
  - e. Save your changes.
7. To assign users to roles, complete the following steps:
  - a. Select the *wsgwauth* project, then edit the **EAR Deployment Descriptor**.
  - b. On the Security tab, select **Gather**. For every security role that you want to assign, repeat the following steps:



- 1) Select a security role.
- 2) Under **WebSphere Bindings**, select the required access level from the following choices:
  - Everyone
  - All authenticated
  - Users/Groups
8. From the pop-up menu for the wsgwauth project, select **Generate Deploy Code**.
9. Export the wsgwauth project as an EAR file.

You are now ready to install the modified copy of the wsgwauth.ear file as described in the final step of Enabling operation-level authorization.

## Invoking Web services over HTTPS

The Web services gateway can invoke Web services that include `https://` in their addresses, if the Java and WebSphere security properties are configured appropriately. This means that one gateway can send a SOAP over HTTPS message direct to another gateway, rather than exporting services and having clients invoke them using HTTPS.

To enable your gateway to send and receive SOAP/HTTPS messages, confirm that your Java and WebSphere security properties are configured as described in the following steps:

1. Check that there is a copy of the `ibmjssse.jar` file in the `install_root/java/jre/lib/ext` directory where `install_root` is the root directory for your installation of IBM WebSphere Application Server (by default `WebSphere/AppServer`).
2. Edit the `install_root/java/jre/lib/security/java.security` security properties file so that it includes entries for both the Sun security provider and the IBM security provider. For example:

```
security.provider.1=sun.security.provider.Sun
security.provider.2=com.ibm.jsse.IBMJSSEProvider
```

The order is significant. The Sun security provider must come before the IBM provider.

3. Use the WebSphere Application Server administrative console to set up the following equivalent system properties:

```
// truststore location
System.setProperty("javax.net.ssl.trustStore",
 "your_truststore_root_directory/TestSSL/key.jks");
// set truststore password
System.setProperty("javax.net.ssl.trustStorePassword",
 "your_truststore_password");
//use ibm reference implementation
System.setProperty("java.protocol.handler.pkgs",
 "com.ibm.net.ssl.internal.www.protocol");
```

## Enabling proxy authentication for the gateway

The gateway requires access to the Internet for invoking Web services and for retrieval of WSDL files. Many enterprise installations use a proxy server in support of Internet routing, and many proxy servers require authentication before they grant access to the Internet.

This requirement is supported in HTTP messaging by a `Proxy-Authorization` message header that contains encoded user name and password credentials.

For messages passing through the gateway, you can enable and disable proxy authentication, and specify whether the authentication credentials are supplied by the service requester or by the gateway. If you specify *requester-supplied* credentials, the credentials in the HTTP message that the gateway receives are re-instantiated by the gateway in the equivalent message that it sends on to the proxy. If you specify *gateway-supplied* credentials, the gateway ignores any credentials in the incoming HTTP message and supplies its own credentials in the equivalent message that it sends on to the proxy.



In certain circumstances, the gateway also creates and sends its own messages (for example for WSDL retrieval). In these cases the gateway always supplies its own credentials to the authenticating proxy. Therefore even if you enable proxy authentication and specify *requester-supplied* credentials, you must still supply credentials for the gateway.

To enable proxy authentication for the gateway, complete the following steps:

1. Display the Web services gateway administrative user interface.
2. In the navigation pane, click the following link:

**Gateway**

- Configure

The gateway configuration form is displayed:

# IBM Web Services Gateway

## Gateway

- **Configure**
- Security
- Back Up
- Restore
- About

## Services

- List
- Deploy
- Remove

## Channels

- List
- Deploy
- Remove

## Filters

- List
- Deploy
- Remove

## Handlers

- List
- Deploy
- Remove

## UDDI References

- List
- Deploy
- Remove

## Configure Gateway

Namespace URI for services

WSDL URI for exported definitions

## Proxy Configuration

Enable proxy authentication

Proxy user

Proxy password

Use Gateway proxy credentials for invoking WebServices

**Note:** You also use the gateway configuration form to set the namespace URI and WSDL URI for the Web services gateway.

3. Enable the **Enable proxy authentication** check box.
4. In the **Proxy user** field, type the proxy user name for the gateway.

**Note:** If you enable proxy authentication then this field is compulsory, even if you also specify *requester-supplied* credentials as described in a subsequent step.

5. In the **Proxy password** field, type the associated proxy password for the gateway.

**Note:** If you enable proxy authentication then this field is compulsory, even if you also specify *requester-supplied* credentials as described in the next step.

6. To set the **Use Gateway proxy credentials for invoking WebServices** check box, complete one of the following two steps:
  - a. To use *requester-supplied* credentials, clear the **Use Gateway proxy credentials** check box.  
With this setting, each incoming message to the gateway from a service requester is expected to contain a valid Proxy-Authorization HTTP message header. This header is re-instantiated by the gateway in the equivalent message that it sends to the proxy.  
For gateway-initiated messaging, such as WSDL retrieval, the gateway supplies its own credentials in the HTTP messages that it sends to the proxy.
  - b. To use *gateway-supplied* credentials, enable the **Use Gateway proxy credentials** check box.  
With this setting, a trust association is established between the gateway and the authenticating proxy. The gateway supplies its own credentials in all messages that it sends to the proxy, and no user name or password is required from service requesters for invoking Web services.
7. Click **Apply Changes**.
8. To provide the application server in which your gateway is running with machine details for the authenticating proxy and for any internal machines that do not require authentication, set system properties in the Java Virtual Machine (JVM) of WebSphere Application Server by completing the following steps:
  - a. Start the WebSphere Application Server administrative server.
  - b. Start the administrative console.
  - c. In the navigation pane, select **Application Servers > your\_server\_name > Process Definition > Java Virtual Machine > Custom Properties**.
  - d. Set the following properties:
    - **http.proxySet** - Set this to true to tell the application server that it is required to work with an authenticating proxy.
    - **http.proxyHost** - Set this to the machine name of the authenticating proxy.
    - **http.proxyPort** - Set this to the port through which the authenticating proxy is accessed. For example 8080
    - **http.nonProxyHosts** - List the internal machines for which authentication is not required for routing through the proxy. Separate each machine name in the list with a vertical bar (|).
    - This list must include the machine on which the gateway is installed.
  - e. Save the settings.
  - f. Stop then restart the application server.
  - g. Close the administrative console.

---

## Tuning the Web services gateway

For information on tuning the application server, see "Tuning performance" in the Information Center.

To tune the gateway, implement as appropriate the performance enhancements described in the following list:

- **Using the file:///filename notation for locally-stored WSDL files**

If you are passing an address for a WSDL file into the gateway, and the WSDL is stored locally, then use the `file:///filename` notation rather than `drive letter:/filename` because `file:///filename` is more performant.

- **Reviewing the gateway storage requirements**

The amount of storage (memory) that the gateway occupies depends upon the number of resources defined (channels, services, filters and so on). The complexity of the WSDL also affects the storage requirements for each Web service. If the resources you deploy to the gateway grow over time, you should review the heap size accordingly.

Although the Java heap storage requirement for each gateway-deployed service is small, if you deploy a large number of services or services with a large number of targets then you can significantly increase the heap space required to run the gateway efficiently. As well as the heap space required for the resource definitions, a large amount of additional heap space is occupied by resource objects instantiated when each service is invoked. The more deployed services or targets that are invoked, the more heap space is required by the gateway for the reusable elements of those services or targets.

The measurements in the following table estimate the Java heap storage space used by the gateway. These measurements were taken for a heap with these settings:

- **-Xms=256**
- **-Xmx=256**

The Web service defined was a simple case with one operation and no complex types.

Resource	Storage space
Java heap requirements for the gateway with the following resources defined: <ul style="list-style-type: none"> <li>• SOAP over HTTP Channel</li> <li>• One Web service</li> <li>• One basic routing filter</li> </ul>	200KB
Java heap requirements for each additional Web service target defined.	1KB
Java heap requirements for each Web service at runtime (that is, those components that are cached for future use and not eligible for garbage collection).	1.2MB for first invocation 100KB for each additional target invoked

These measurements are provided for illustrative purposes only. To review the gateway storage requirements for your environment, you should take your own measurements for that environment according to the formula described in the next paragraph.

For each Web service defined to the gateway, and for each service subsequently invoked, the Java heap working set is increased. The increase in heap occupancy is a function of base gateway requirements, plus the number of Web services defined, plus the number of unique Web services invoked during the lifetime of the gateway. This can be expressed as:

$$gateway\_server\_storage + n*def + rdef1 + n2*rdef2$$

where:

- *gateway\_server\_storage* is the start-up footprint of the gateway server (not measured in the preceding table), plus any basic deployments of channels, filters and so on (200KB in the preceding table). This startup storage requirement varies depending upon factors such as the initial configuration of the server, and whether or not any other applications are running.
- *n* is the number of Web services (or targets) defined.
- *def* is the number of bytes used for each definition.
- *rdef1* is the number of bytes used for the first invocation of the Web service.
- *n2* is the number of unique Web services (or targets) executed during the lifetime of the gateway.
- *rdef2* is the number of bytes used for each additional invocation of a unique Web service target (that is, those objects for each target that are cached for future use and not eligible for garbage collection).

• **5.1+ Setting application class-loader policy to enable local EJB calls**

Every request that arrives at a channel is passed by the channel to the gateway. This communication is made through either a remote or a local EJB call. If you set the application class-loader policy for the application server to MULTIPLE, then this forces a remote EJB call and the request has to be serialized then de-serialized. If you set the application class-loader policy to SINGLE, then the channel makes a local call and passes the request by reference.

If you set the application class-loader policy to SINGLE, then the gateway can also call a filter through a local EJB call, provided that the filter has been written to provide the local interfaces.

For information on how to write a filter that can be called through a local EJB call, see [Writing a filter for the Web services gateway](#).

- **5.1+** **Enabling selective SOAP parsing**

When you enable selective SOAP parsing for a specific Web service, the gateway only parses the headers of an incoming message for that service. The message body is passed through unchecked. Enable this option, as described in [Deploying Web services to the Web services gateway](#), if you need to optimize the speed at which messages pass through the gateway. For selective SOAP parsing to work, the application class-loader policy for the application server must be set to SINGLE.

For more information, see [Selective SOAP parsing](#).

## Selective SOAP parsing

By default, the gateway parses the whole of every SOAP message that it receives. However, you can speed up the passage of messages through the gateway by choosing selective SOAP parsing.

When you enable selective SOAP parsing for a specific Web service, the gateway only parses the headers of incoming messages for that service. The message body is passed through unchecked.

You enable this option if you need to optimize the speed at which messages pass through the gateway, or if SOAP parsing is already being carried out by another process before the message reaches the gateway (because, with selective SOAP parsing enabled, you can send in a SOAP message that contains an illegal body and it will pass through the gateway without stopping the process or causing an error to be reported).

Even with selective SOAP parsing enabled, you can still use gateway filters that act upon the message body (provided that the message body is legal), because the body is still passing through the gateway even though it is not being parsed.

You enable or disable the selective SOAP parsing option, for each deployed gateway service, as described in [Deploying Web services to the Web services gateway](#).

For selective SOAP parsing to work, you must set the application class-loader policy to SINGLE. For more information, see [Tuning the Web services gateway](#).

Selective SOAP parsing is only supported by the SOAP over HTTP channel. If incoming messages are received on any other channel, they are fully parsed whether or not you have enabled selective SOAP parsing.

---

## Troubleshooting the Web services gateway

This topic provides hints to help you resolve problems you experience when using the Web services gateway.

For information on resolving WebSphere-level problems, see ["Diagnosing and fixing problems"](#) in the Information Center.

For current information available from IBM Support on known problems and their resolution, see the [IBM Support page](#).

IBM Support has documents that can save you time gathering information needed to resolve this problem. Before opening a PMR, see the IBM Support ["must gather"](#) page.

If you report a gateway-related problem to IBM Service and Support, include the gateway release and build information that is listed under the **About** option of the Web services gateway administrative user interface.

To help you identify and resolve gateway-related problems, use the WebSphere Application Server trace and logging trace and logging facilities. (See "Enabling trace" and "Setting up component trace (CTRACE)" in the Information center.) To enable trace for the gateway, set the application server trace string to `com.ibm.wsgw.*=all=enabled`. If you encounter a problem that you think might be related to the gateway, you can check for error messages in the WebSphere Application Server administrative console, and in the application server `stdout.log` file. You can also enable the application server debug trace to provide a detailed exception dump.

The gateway administrative user interface uses cascading style sheets to lay out its pages, and Javascript to monitor progress and to advise you as you fill in each on-screen form. Your Web browser must therefore support Javascript and cascading style sheets, and it must be configured so that Javascript and style sheets are enabled. This configuration depends on which browser you use. For example for Netscape, you select **Edit > Preferences**, click Advanced in the Category pane, then confirm that the **Enable Javascript** and **Enable style sheets** check boxes are selected.

A list of the gateway run-time system messages, with details of what each message means, is provided in Web services gateway messages.

Here is a checklist of common problems:

- **You are about to install an upgraded gateway, but you do not want to lose your existing gateway configuration.**

When you complete the installation of an upgraded gateway, any previously configured gateway is replaced with an upgraded but empty gateway. To preserve an existing gateway configuration, you need to save the configuration before you upgrade the gateway, then restore the configuration after the upgrade is installed. For detailed instructions see Preserving an existing gateway configuration.

- **You have managed to deploy your Web service in the Web services gateway but you are getting a class cast exception when you invoke the operation which takes an integer parameter.**

Check that your client is using the version of the `soap.jar` file that is supplied in the WebSphere Application Server `/AppServer/lib/app` directory. If you enable trace, you can look in the trace for the request `<SOAP-ENV:Envelope xmlns:SOAP-ENV="http://schemas.xmlsoap.org/soap/envelope/" xmlns:xsi="http://www.w3.org/1999/XMLSchema-instance" xmlns:xsd="http://www.w3.org/1999/XMLSchema">`

The gateway expects the 2001 version of the XML schema. Older versions of the `soap.jar` file (including 2.2) generate 1999 schema. If you have the `soap.jar` file that is supplied with WebSphere Application Server in the client class path, you can see 2001 schema in the request.

- **The persistent state of the Web services gateway has become out of sync with the channel applications.**

This mismatch can occur if you remove and reinstall the Apache SOAP channel applications. If this operation is necessary, then either verify that all corresponding channels configured with the Web services gateway are removed, or remove and reinstall the Web services gateway at the same time.

**Note:** The Web services gateway application (`wsgw.ear`) must be installed before the channel and filter applications. If the gateway application needs to be reinstalled, all channels and filters must be uninstalled first, then reinstalled after the gateway application.

- **5.1+ You get listener port time out errors when large messages are passed using the synchronous SOAP over JMS channel.**

As with any synchronous channel, time out errors can occur. To minimize the frequency of time out errors, increase the time out settings for the channel. If the problem persists, then disable trace and logging trace and logging for the gateway (by setting the application server trace string to `com.ibm.wsgw.*=all=disabled`).

- **You are getting SOAP fault messages, but cannot determine the precise problem from the fault message.**

If you receive a SOAP fault message with a faultstring that is just the value of one of the parameters of the invocation, that means the parameter value is not valid. For example if you have a service that expects an int parameter and you send it a message containing the value “1.1”, then the fault message you receive contains 1.1 as the fault string:

```
<faultcode>SOAP-ENV:Client</faultcode>
<faultstring>1.1</faultstring>
```

This message is consistent with Apache SOAP behavior, and is not correctable by the gateway.

If you receive a SOAP fault message containing an element that is not present in the WSDL for the target service, then the error message thrown can be difficult to identify. There are two possible scenarios:

- The WSDL is deployed to use **Generic Classes**. In this case the returned SOAP message contains an `IllegalArgumentException` exception. For example:

```
[Attributes={}] [faultCode=SOAP-ENV:Server]
[faultString=com.ibm.wsgw.WSGWException:
WSGW0043E: Exception while executing
 operation createEntry service ExchangeService.
Exception: org.apache.wsif.WSIFException: SOAPException:
SOAP-ENV:ClientNo mapping found for 'com.ibm.jrom.JROMValue'
using encoding style 'http://schemas.xmlsoap.org/soap/encoding/';
nested exception is:
 [SOAPException: faultCode=SOAP-ENV:Client; msg=No mapping found for
'com.ibm.jrom.JROMValue' using encoding style
'http://schemas.xmlsoap.org/soap/encoding/';
targetException=java.lang.IllegalArgumentException:
No mapping found for 'com.ibm.jrom.JROMValue' using encoding style
'http://schemas.xmlsoap.org/soap/encoding/'.]]
[faultActorURI=/wsgwsoap1/soaprprouter]
...
```

- The WSDL is deployed to use **Deployed Classes**. In this case an empty message is returned. For example:

```
[Attributes={}] [faultCode=null] [faultString=null] [faultActorURI=null]
[DetailEntries=] [FaultEntries=]
```

These messages are consistent with Apache SOAP behavior, and are not correctable by the gateway.

- **When you are specifying the security settings that are applied between the service requester (the client) and the gateway, you set Identification to Digital Signature. When you click Apply Changes, the following “AuthMethod signature is not valid” error message is displayed:**

```
WSGW0009E: Failed to deploy service. Exception: Error change service properties:
com.ibm.wsspi.wssecurity.SoaSecurityException: WSEC0105E: LoginConfig: AuthMethod Signature is not valid.
```

This problem only occurs when you set **Identification** to **Digital Signature**. The problem is caused by a mismatch between the request bindings and the response bindings. For example, you might have set **Request Sender** signing information on the request bindings instead of **Request Receiver**, or **Response Receiver** on the response bindings instead of **Response Sender**. To fix the problem, set matching request and response bindings. For instructions on setting these parameters, see

- Request receiver binding collection
- Response sender binding collection
- Request sender binding collection
- Response receiver binding collection



- **When you run secured Web services through a Web services gateway using LTPA authentication with a certificate store defined, an “Unable to build a valid CertPath” error similar to the following is displayed:**

```
[10/23/03 12:31:40:393 CDT] 78a54799 enterprise I
com.ibm.ws.webservices.engine.enterprise
TRAS0014I: The following exception was logged WebServicesFault
faultCode: {http://schemas.xmlsoap.org/ws/2003/06/secext}InvalidSecurityToken
faultString: WSEC5085E: Unable to build a valid CertPath:
java.security.cert.CertPathBuilderException:
PKIXCertPathBuilderImpl could not build a valid CertPath.;
internal cause is: java.security.cert.CertPathValidatorException:
The certificate issued by CN=Int CA2, OU=TRL, O=IBM, ST=Kanagawa, C=JP is not trusted;
internal cause is: java.security.cert.CertPathValidatorException: Certificate chaining error
```

This problem occurs when one of the defined certificate stores on the Web services gateway administrative client is missing the certificate path.

To correct this problem,

1. Add the certificate path to the certificate store configuration. For example, if you are using the sample certificate store from WebSphere Application Server, the path is:

```
X509 Certificate Path: /usr/WebSphere/AppServer/etc/ws-security/samples/intca2.cer
```

2. Run the client again.

- **When you run secured Web services through the Web services gateway, a “Failed authentication” error similar to the following is displayed:**

```
[10/3/03 21:35:48:597 CDT] 33c558 enterprise I
com.ibm.ws.webservices.engine.enterprise TRAS0014I: The following exception was
logged WebServicesFault
faultCode: {http://schemas.xmlsoap.org/ws/2003/06/secext}FailedAuthentication
faultString: WSEC5078E: Login failed: javax.security.auth.login.LoginException:
java.lang.NullPointerException
at com.ibm.wsspi.wsssecurity.auth.module.WSSecurityMappingModule._login(WSSecur
ityMappingModule.java:152)
```

The problem occurs when the Web services gateway does not enable global security and there is no LDAP user registry information to authenticate the incoming user login.

To correct this problem, enable global security on the Web services gateway application server using the same LDAP user registry that the deployment manager uses. See Enabling gateway-level authentication for more information.

**Note:** Before enabling global security on the Web services gateway, you need to uninstall the Web services gateway channels and applications, then re-install them after security is enabled.

- **When you run secured Web services through the Web services gateway, a “Failed login” error similar to the following is displayed:**

```
WSEC5078E: Login failed: com.ibm.websphere.security.auth.WSLoginFailedException:
AuthenLoginModule: Authentication failed reason = 2
```

This problem occurs when single sign-on is enabled and the same password is used on both the deployment manager and the Web services gateway application server. The login fails even though the deployment manager and the Web services gateway application server are both enabled with security, with the same LDAP user registry and with the same LTPA password. The problem occurs because the key that is generated and passed between the services is randomly generated, and is therefore different for each server even though the same password is used.

To correct this problem, export the key from one cell then import it into another.

- **You are enabling operation-level authorization, but when you install wsgwauth.ear, an error message appears in the WebSphere Application Server administrative console detailing a Java Naming and Directory Interface (JNDI) problem.**

Check that you entered, in the “EJB References” for the authorization session bean, the correct JNDI name of the imported Web service enterprise bean. Note that this home is case sensitive.

- 5.1+** You get JNDI lookup errors when you use the same names for JMS messaging queues and queue connection factories that run on application servers on different machines. You should not use the same names for messaging queues and queue connection factories that run on application servers on different machines, because the Web Services Invocation Framework (WSIF) (which the gateway uses for service invocation) always looks first for JMS destinations locally, and only uses the full JNDI reference if it cannot find the destination locally. If you deploy to the gateway a Web service that is hosted on a remote machine, and the same names are used for messaging queues and queue connection factories on the remote machine and on the machine on which the gateway is hosted, then WSIF finds and uses the local queues even if the remote JNDI destination is provided in full in the WSDL service definition.
  - You are trying to have a Web services gateway send a SOAP/HTTPS message to another Web services gateway, and you are receiving a MalformedURLException error.**

The Web services gateway can invoke Web services that include `https://` in their addresses, if the Java and WebSphere security properties are configured accordingly. To check your security property settings, see the topic Invoking Web services over HTTPS
  - You clear 'Authorization Policy - Control access to this service' from the deployment details for a Web service, and you find the service no longer works.**

A number of tasks are required to disable security. Clearing the check box for 'Authorization Policy - Control access to this service' still leaves WebSphere Application Server security in place, so basic authentication might still be required.

To disable security fully, use the WebSphere Application Server administrative console Security Center to disable Global Security.
  - You experience problems with handling Document style SOAP messages.**

If you are using the Apache SOAP channel, then the SOAP message format must be RPC style. To handle Document style SOAP messages, use the SOAP over HTTP channel (which supports both RPC style and Document style SOAP messages).
  - You experience problems with handling SOAP messages with attachments.**

To handle SOAP messages with attachments, use the SOAP over HTTP channel.
  - You experience problems with a Web service that uses complex data types. The same service works fine, when not using the gateway.**

You need to do one of two things to support Web services that use complex types in the Web services gateway:

    - Set the message part representation for the service to **Generic classes**.
    - Set the message part representation for the service to **Deployed Java classes**, and make the original classes available to the application server (either by updating the class path or by putting the JAR file somewhere like `install_root/lib/app`, where `install_root` is the root directory for your installation of IBM WebSphere Application Server). If you do not have the original classes, you can use `wsdl2java` to generate Java applications that can contain values of the complex type for the service, then compile the beans into a JAR file and make this file available to the application server.

For more information on the factors to consider when choosing between these options, see Data type representation - Choosing between Generic classes and Deployed Java classes.
- Note:**
- If your Web service has non-bean parameters (that is, it requires a custom serializer and deserializer) then it is not supported by the current release of the gateway.
  - Not all complex types that are expressible in XML schema are supported by the current release of the gateway. For more information, see Web services gateway - Supported types.
  - See also Complex data types - mapping namespaces to packages.
- You develop a filter as described in Writing a filter for the Web services gateway. When you deploy the filter to the gateway, the following error message appears: WSGW0035E: Filter context version 5 not supported.** You are trying to deploy your filter to a server other than WebSphere Application Server Enterprise Version 5 or WebSphere Business Integration Server Foundation Version

5.1. Your filter uses the gateway WorkArea to maintain context, and the gateway WorkArea is only available with IBM WebSphere Application Server Enterprise Version 5 or WebSphere Business Integration Server Foundation Version 5.1.

## Web services gateway messages

This topic contains a list of the Web services gateway run-time system messages, with details of what each message means.

WebSphere system messages are logged from a variety of sources, including application server components and applications. Messages logged by application server components and associated IBM products start with a unique message identifier that indicates the component or application that issued the message.

For more information about the message identifier format, see the topic "Message reference" in the Information Center.

**GWIN0003I: Looking for {childtype} {childname} under {parenttype} {parentname}. Values are: {0}**

**Explanation:** Gateway installation message from file setupwsgw.jac].

**User Response:**

**GWIN0004I: Looking for {childtype} {childname} under {parenttype} {parentname} and {parenttype2} {parentname2}. Values are: {0}**

**Explanation:** Gateway installation message from file setupwsgw.jac].

**User Response:**

**GWIN0005I Conflict found with existing {childtype} {childname}. Values are: {0}**

**Explanation:** Gateway installation message from file setupwsgw.jac].

**User Response:**

**GWIN0012I: Not creating requested {childtype}. Value is: {0}**

**Explanation:** Gateway installation message from file setupwsgw.jac].

**User Response:**

**GWIN0006I: Seeking {parenttype} with requested id of {parentname}. Values are: {0}**

**Explanation:** Gateway installation message from file setupwsgw.jac].

**User Response:**

**GWIN0007I: Seeking {parenttype} with requested id of {parentname} under {parenttype2} {parentname2}. Values are: {0}**

**Explanation:** Gateway installation message from file setupwsgw.jac].

**User Response:**

**GWIN0008I: Attempting to create {childtype} under {parenttype} of {parentID}. Values are: {0}**

**Explanation:** Gateway installation message from file setupwsgw.jac].

**User Response:**

**GWIN0009I: Create command that will be issued is: {0}**

**Explanation:** Gateway installation message from file setupwsgw.jac].

**User Response:**

**GWIN0010I: {childtype} {childId} was successfully created. Values are: {0}**

**Explanation:** Gateway installation message from file setupwsgw.jac].

**User Response:**

**GWIN0101I: Attempting to save new configuration.**

**Explanation:** Gateway installation message from file setupwsgw.jac].

**User Response:**

**GWIN0102I: Changes saved successfully.**

**Explanation:** Gateway installation message from file setupwsgw.jac].

**User Response:**

**GWIN0103I: Changes were not saved on this call.**  
**Explanation:** Gateway installation message from file setupwsgw.jac1.  
**User Response:**

**GWIN0014I: Looking for builtin\_rra.**  
**Explanation:** Gateway installation message from file setupwsgw.jac1.  
**User Response:**

**GWIN0015I: List for J2CResourceAdapter returned {0} members. Value is: {1}**  
**Explanation:** Gateway installation message from file setupwsgw.jac1.  
**User Response:**

**GWIN0049I: List for Applications returned {0} members. Value is: {1}**  
**Explanation:** Gateway installation message from file setupwsgw.jac1.  
**User Response:**

**GWIN0016I: Hunting J2CResourceAdapter associated with Node {nodename}. Value is: {0}**  
**Explanation:** Gateway installation message from file setupwsgw.jac1.  
**User Response:**

**GWIN0017I: Using {rraID} as builtin\_rra. Value is: {0}**  
**Explanation:** Gateway installation message from file setupwsgw.jac1.  
**User Response:**

**GWIN0018I: Using provider class of {implclass} with a classpath of {clpath}. Values are: {0}**  
**Explanation:** Gateway installation message from file setupwsgw.jac1.  
**User Response:**

**GWIN0019I: Installing to server {servername}, node {nodename} using database type of {dbtype}. Values are: {0}**  
**Explanation:** Gateway installation message from file setupwsgw.jac1. The database parameter is not used in this version of the gateway.  
**User Response:**

**GWIN0023I: Checking for installed WSGW of name {appname}. Value is: {0}**  
**Explanation:** Gateway installation message from file setupwsgw.jac1.  
**User Response:**

**GWIN0025I: ApplicationManager not running, so application will not need to be stopped.**  
**Explanation:** Gateway installation message from file setupwsgw.jac1.  
**User Response:**

**GWIN0026I: Stopping application of name {0}. Value is: {1}**  
**Explanation:** Gateway installation message from file setupwsgw.jac1.  
**User Response:**

**GWIN0027W: stopApplication command for application {0} caught exception Exc. Application might not have been running on this server. Values are: {1}**  
**Explanation:** Gateway installation message from file setupwsgw.jac1.  
**User Response:**

**GWIN0028I: Application {0} stopped successfully. Value is: {1}**  
**Explanation:** Gateway installation message from file setupwsgw.jac1.  
**User Response:**

**GWIN0029I: Removing application {0}. Value is: {1}**  
**Explanation:** Gateway installation message from file setupwsgw.jac1.  
**User Response:**

**GWIN0030I: Application {0} removed successfully. Value is: {1}**  
**Explanation:** Gateway installation message from file setupwsgw.jac1.  
**User Response:**

**GWIN0050I: Install of WSGW application {0} successful. Value is: {0}**  
Explanation: Gateway installation message from file setupwsgw.jac].  
User Response:

**GWIN0051I: Application Manager {0} found. Value is: {1}**  
Explanation: Gateway installation message from file setupwsgw.jac].  
User Response:

**GWIN0031I: Adding resource bundles to repository.**  
Explanation: Gateway installation message from file setupwsgw.jac].  
User Response:

**GWIN0035I: Url prefix found. Updating it to discoveryURL. Value is: {0}**  
Explanation: Gateway installation message from file setupwsgw.jac].  
User Response:

**GWIN0038I: ws.ext.dir processing starting.**  
Explanation: Gateway installation message from file setupwsgw.jac].  
User Response:

**GWIN0013I: JVM is: {0}**  
Explanation: Gateway installation message from file setupwsgw.jac].  
User Response:

**GWIN0040I: Out of {0} properties we located {1} matches at positions {poslist}. Values are: {2}**  
Explanation: Gateway installation message from file setupwsgw.jac].  
User Response:

**GWIN0042I: SYSPROP is: {0}**  
Explanation: Gateway installation message from file setupwsgw.jac].  
User Response:

**GWIN0043I: ws.ext.dir has been set with new sysprop. Value is: {0}**  
Explanation: Gateway installation message from file setupwsgw.jac].  
User Response:

**GWIN0107I: Changes saved successfully.**  
Explanation: Gateway installation message from file setupwsgw.jac].  
User Response:

**GWIN0044I: ws.ext.dir update skipped, required changes already present.**  
Explanation: Gateway installation message from file setupwsgw.jac].  
User Response:

**GWIN0045I: ws.ext.dir processing step complete.**  
Explanation: Gateway installation message from file setupwsgw.jac].  
User Response:

**GWIN0047I: Issuing nodeSync.**  
Explanation: Gateway installation message from file setupwsgw.jac].  
User Response:

**GWIN0048I: WSGW successfully installed.**  
Explanation: Gateway installation message from file setupwsgw.jac].  
User Response:

**GWIN0052I: WSGW successfully installed. Please restart server {0} to activate configuration changes. Value is: {1}**  
Explanation: Gateway installation message from file setupwsgw.jac].  
User Response:

**GWIN0053I: WSGWResourceBundles.jar file was not added to the repository - might already exist.**  
Explanation: Gateway installation message from file setupwsgw.jac].

**User Response:**  
**GWIN0054I: Building new props.**  
**Explanation:** Gateway installation message from file setupwsgw.jacl.

**User Response:**  
**GWIN6001E: Incorrect number of arguments passed to script.**  
**Explanation:** Gateway installation message from file setupwsgw.jacl.

**User Response:**  
**GWIN6002E: Usage: wsadmin -f setupwsgw.jacl <deployMgrPath> <serverName> <nodeName> <dbName> <dbuserid> <db2password> <db2zipLocation>**  
**Explanation:** Gateway installation message from file setupwsgw.jacl. The database parameters are not used in this version of the gateway.

**User Response:**  
**GWIN6003E: <db2userid> <db2password> <db2zipLocation> only required for db2 jdbc provider.**  
**Explanation:** Gateway installation message from file setupwsgw.jacl. Not used in this version of the gateway.

**User Response:**  
**GWIN6005E: Removal of {childtype} {childname} caught exception Exc. Values are: {0}**  
**Explanation:** Gateway installation message from file setupwsgw.jacl.

**User Response:**  
**GWIN6006E: An exception Exc occurred while creating {childtype}. Values are: {0}**  
**Explanation:** Gateway installation message from file setupwsgw.jacl.

**User Response:**  
**GWIN6008E: Unable to find requested {parenttype} of {parentname}. Values are: {0}**  
**Explanation:** Gateway installation message from file setupwsgw.jacl.

**User Response:**  
**GWIN6009E: List command for J2CResourceAdapter caught exception Exc. Value is: {0}**  
**Explanation:** Gateway installation message from file setupwsgw.jacl.

**User Response:**  
**GWIN6010E: No J2CResourceAdapter objects available.**  
**Explanation:** Gateway installation message from file setupwsgw.jacl.

**User Response:**  
**GWIN6012E: Uninstall of application {0} caught exception Exc. Values are: {1}**  
**Explanation:** Gateway installation message from file setupwsgw.jacl.

**User Response:**  
**GWIN6013E: Install of WSGW application caught exception Exc. Value is: {0}**  
**Explanation:** Gateway installation message from file setupwsgw.jacl.

**User Response:**  
**GWIN6014E: Could not get JVM.**  
**Explanation:** Gateway installation message from file setupwsgw.jacl.

**User Response:**  
**GWIN6017E: Cannot find nodeSync MBean.**  
**Explanation:** Gateway installation message from file setupwsgw.jacl.

**User Response:**  
**GWIN6018E: nodeSync failed. WSGW applications may not be fully installed.**  
**Explanation:** Gateway installation message from file setupwsgw.jacl.

**User Response:**  
**GWIN6019E: Install of WSGW application {0} caught exception Exc. Values are: {1}**  
**Explanation:** Gateway installation message from file setupwsgw.jacl.



**User Response:**

**GWIN6020E: <dbName> only required for DB2/Cloudscape.**

**Explanation:** Gateway installation message from file setupwsgw.jacl. Not used in this version of the gateway.

**User Response:**

**GWIN6021E: Use all forward slashes (/) to avoid problems with escaping back slashes (\\).**

**Explanation:** Gateway installation message from file setupwsgw.jacl.

**User Response:**

**GWIN6022E: Example: wsadmin -f setupwsgw.jacl C:/WebSphere/DeploymentManager  
bunny\_server1 bunny WSGWDB db2admin db2passwd C:/SQLLIB/java/db2java.zip**

**Explanation:** Gateway installation message from file setupwsgw.jacl. The database parameters are not used in this version of the gateway.

**User Response:**

**GWIN6101E: Error saving configuration, changes not saved due to exception Exc. Value is: {0}**

**Explanation:** Gateway installation message from file setupwsgw.jacl.

**User Response:**

**WSGW0001E: Channel name {0} from gateway configuration differs from that in JNDI: {1}**

**Explanation:** The name specified for the channel does not match the name of the channel as defined within the EAR file.

**User Response:** Ensure that the channel name is specified correctly

**WSGW0002E: Error storing endpoint address. Exception: {0}**

**Explanation:** An unexpected exception occurred when storing the endpoint address for a channel.

**User Response:** Contact IBM Support

**WSGW0003E: Error retrieving endpoint address. Exception: {0}**

**Explanation:** An unexpected exception occurred when retrieving the endpoint address for a channel.

**User Response:** Contact IBM Support

**WSGW0004E: Not used**

**Explanation:**

**User Response:**

**WSGW0005E: Error retrieving channel name. Exception: {0}**

**Explanation:** An unexpected exception occurred when retrieving the channel name from JNDI.

**User Response:** Contact IBM Support

**WSGW0006E: Error deploying service to {1}. Exception: {0}**

**Explanation:** An unexpected error occurred deploying the service to the given component.

**User Response:** This error may be caused by a previous failure. Try redeploying the service using a different gateway service name. If that fails, reinstalling the channel and gateway applications may remove the problem.

**WSGW0007E: Error getting endpoint URL from channel {0}. Exception: {1}**

**Explanation:** An unexpected error occurred generating the endpoint URL for the given channel.

**User Response:** Contact IBM Support

**WSGW0008E: Could not determine default port name for target service {0}**

**Explanation:** There are no ports in the WSDL defined for the target service that are supported by currently available WSIF providers or there is an error in the WSDL file associated with the port definition or a namespace it uses.

**User Response:** Either ensure that a WSIF provider is correctly configured for the port in the WSDL, or ensure that the WSDL contains correctly specified port information.

**WSGW0009E: Failed to deploy service. Exception: {0}**

**Explanation:** An unexpected error occurred trying to deploy the service.



- User Response:** Contact IBM Support
- WSGW0010E: The namespaceURI attribute cannot be changed when there are active services**  
**Explanation:** The namespaceURI is used to generate WSDL for gateway services. If this global setting is changed then current WSDL becomes invalid.
- User Response:** Either remove all channels or all gateway services from the gateway configuration and retry the change.
- WSGW0011E: Not used**  
**Explanation:**
- User Response:**
- WSGW0012E: Not used**  
**Explanation:**
- User Response:**
- WSGW0013E: Could not locate home {0}. Exception: {1}**  
**Explanation:** The specified home location could not be found in JNDI.
- User Response:** Ensure that the home location is specified correctly, and that it appears in JNDI.
- WSGW0014W: Could not access local home {0}. Attempting remote home {1}.**  
**Explanation:** The local home was found in JNDI but could not be used as the application server “Application classloader policy” is set to “Multiple”. The remote home will be used for this component.
- User Response:** Ensure that the application server “Application classloader policy” is set to “Single” to allow the gateway to use local interfaces. Ignore this message if you require “Application classloader policy” setting to be “Multiple”.
- WSGW0015E: Could not create instance of class {0}. Exception: {1}**  
**Explanation:** The gateway failed to create an instance of the specified Java class.
- User Response:** Ensure that the Java class has a public constructor with no parameters.
- WSGW0016E: Could not locate class {0}. Exception: {1}**  
**Explanation:** The gateway failed to locate the specified Java class.
- User Response:** Ensure that the Java class is visible to the gateway application’s classloader.
- WSGW0017E: Not used**  
**Explanation:**
- User Response:**
- WSGW0018E: Not used**  
**Explanation:**
- User Response:**
- WSGW0019E: Failed to clone definition. Exception: {1}**  
**Explanation:** An unexpected error occurred cloning a WSDL definition.
- User Response:** Contact IBM Support
- WSGW0020E: Error while loading mapped type class {0}. Exception: {1}**  
**Explanation:** An error occurred while trying to load the given Java class which represents a type in the deployed WSDL for a target service.
- User Response:** Ensure that the Java class is visible to the gateway application’s classloader.
- WSGW0021E: Expected WSDL definition to contain a <wsdl:type> element with a schema from one of the '{0}', '{1}', or '{2}' namespaces**  
**Explanation:** Schema types in WSDL definitions must be declared using one of the specified XML schema namespaces.
- User Response:** Update the WSDL definition to use the appropriate namespace.
- WSGW0022E: Unexpected Schema->Java problem when parsing WSDL file. Exception: {0}**  
**Explanation:** An unexpected exception occurred when parsing a WSDL file. This may be due to unsupported elements in the WSDL.

**User Response:** Contact IBM Support

**WSGW0023E: Unexpected Schema->JROM problem when parsing WSDL file. Exception: {0}**

**Explanation:** An unexpected exception occurred when parsing a WSDL file. This may be due to unsupported elements in the WSDL.

**User Response:** Contact IBM Support

**WSGW0024E: Channel {0} cannot be removed because it is being used by a deployed service**

**Explanation:** channels can only be removed when they are not in use by gateway services.

**User Response:** Remove the channel from gateway services to which it is deployed before removing the channel.

**WSGW0025E: Target service identity cannot be specified as null**

**Explanation:** A target service can only be selected using a non-null valid for the identity.

**User Response:** Modify the calling code to ensure that the target service identity value is never null.

**WSGW0026E: Invalid gateway service name {0}. The name must be a valid XML schema NCNAME.**

**Explanation:** The name specified for the gateway service does not conform to the required definition.

**User Response:** Correct the gateway service name so that it is a valid XML schema NCNAME.

**WSGW0027E: Port {0} does not exist for target service {1}**

**Explanation:** The requested port does not exist for the target service.

**User Response:** Ensure that a valid port is requested, or update the target service WSDL to contain a port of the requested name.

**WSGW0028E: No binding for port {0} for target service {1}**

**Explanation:** The requested port for the target service does not have a binding defined in the WSDL definition of the service.

**User Response:** Ensure that the target service WSDL has a binding for the requested port, or use a different port name.

**WSGW0029E: No portType for binding {0} for port {1} for target service {2}**

**Explanation:** The requested port for the target service does not have a portType defined in the WSDL definition of the service.

**User Response:** Ensure that the target service WSDL has a portType for the requested port, or use a different port name.

**WSGW0030E: Not used**

**Explanation:**

**User Response:**

**WSGW0031E: Channel name {0} already exists**

**Explanation:** The name specified for the channel is the same as that of a channel that is currently deployed.

**User Response:** Choose a different name for the channel, or remove the existing channel of the given name.

**WSGW0032E: Channel name {0} not found**

**Explanation:** No channel is currently deployed with the given name.

**User Response:** Use the name of a channel that is currently deployed.

**WSGW0033E: Filter {0} cannot be removed because it is being used by a deployed service**

**Explanation:** Filters can only be removed when they are not in use by gateway services.

**User Response:** Remove the filter from gateway services to which it is deployed before removing the filter.

**WSGW0034W: Invocation of filter {0} failed. Exception: {1}**

**Explanation:** An unexpected exception was thrown during processing of the given filter.

**User Response:** Contact IBM Support

**WSGW0035E: Filter context version {0} not supported**

**Explanation:** The context version that the filter requires is not supported by this version of the gateway.

**User Response:** Ensure that the filter is requesting the correct context version. It may be necessary to upgrade the gateway to support the filter.

**WSGW0036E: Target service identity information {0} not matched for gateway service {1}**

**Explanation:** A target service was requested by identity, but the identity information does not match any currently deployed target service.

**User Response:** Ensure that the identity information is correct, and that there is a target service deployed to the given gateway service with the right identity information.

**WSGW0037E: Filter name {0} already exists**

**Explanation:** The name specified for the filter is the same as that of a filter that is currently deployed.

**User Response:** Choose a different name for the filter, or remove the existing filter of the given name.

**WSGW0038E: Filter name {0} not found**

**Explanation:** No filter is currently deployed with the given name.

**User Response:** Use the name of a filter that is currently deployed.

**WSGW0039E: Error loading state from {0}. Exception {1}**

**Explanation:** An unexpected exception occurred loading the state of the gateway from the given location.

**User Response:** Ensure that the given location is visible to the gateway application.

**WSGW0040E: Failed to convert definition to string. Exception: {0}**

**Explanation:** An unexpected exception occurred converting a WSDL definition into a string in order to display it or return it to an application.

**User Response:** Contact IBM Support

**WSGW0041E: Failed to save state. Exception {0}**

**Explanation:** An unexpected exception occurred when saving the state of the gateway.

**User Response:** Contact IBM Support

**WSGW0042W: No target services available to get service definition**

**Explanation:** A request was made for the WSDL definition for the gateway service, however no target services have been defined for the gateway service, so it is not possible to generate a WSDL definition.

**User Response:** Deploy one or more target services to the gateway service.

**WSGW0043E: Exception while executing operation {0} service {1}. Exception: {2}**

**Explanation:** An unexpected exception occurred when passing a request on to a target Web service.

**User Response:** Ensure that the gateway service and target service are correctly deployed (using the correct message part representation). Ensure that the target service is available and responds correctly to direct requests (i.e. not through the gateway).

**WSGW0044E: Filter position {0} invalid**

**Explanation:** The specified position for addition or removal of the filter was not valid.

**User Response:** Ensure a valid value is specified. The value should be -1, 0 or a positive integer.

**WSGW0045E: Filter not found in list**

**Explanation:** An attempt was made to remove a filter from a gateway service specifying -1 as the index, but the filter is not in the list at all.

**User Response:** Ensure that the correct filter is specified.

**WSGW0046E: Channel {0} already defined for gateway service {1}**

**Explanation:** The given channel has already been defined for the gateway service.

**User Response:** Ensure that the correct channel name is specified.

**WSGW0047E: Channel {0} not defined for gateway service {1}**

**Explanation:** The channel cannot be removed from the gateway service as it is not currently defined for the gateway service.

**User Response:** Ensure that the correct channel name is specified.

**WSGW0048E: UDDI reference {0} already defined for gateway service {1}**

**Explanation:** The given UDDI reference has already been defined for the gateway service.

**User Response:** Ensure that the correct UDDI reference name is specified.

**WSGW0049E: UDDI reference {0} not defined for gateway service {1}**

**Explanation:** The UDDI reference cannot be removed from the gateway service as it is not currently defined for the gateway service.

**User Response:** Ensure that the correct UDDI reference name is specified.

**WSGW0050E: Target service with location {0} already defined for gateway service {1}**

**Explanation:** The given target service location has already been defined for the gateway service.

**User Response:** Ensure that the correct target service location is specified.

**WSGW0051E: Target service with location {0} not defined for gateway service {1}**

**Explanation:** The target service location cannot be removed from the gateway service as it is not currently defined for the gateway service.

**User Response:** Ensure that the correct target service location is specified.

**WSGW0052E: Target service with location {0} was not found for gateway service {1}**

**Explanation:** The target service WSDL definition could not be obtained from the given location.

**User Response:** Ensure that the correct target service location is specified.

**WSGW0053E: gateway service {0} cannot be removed as active entities and force not specified**

**Explanation:** A gateway service with one or more target services, channels, filters or UDDI references cannot be removed.

**User Response:** Remove the target services, channels, filters and UDDI references from the gateway service.

**WSGW0054E: An exported definition for gateway service {0} is not available as there are no defined channels for the service**

**Explanation:** A request was made for the WSDL definition for the gateway service, however no channels have been defined for the gateway service, so it is not possible to generate a WSDL definition.

**User Response:** Deploy one or more channels to the gateway service.

**WSGW0055E: Not used**

**Explanation:**

**User Response:**

**WSGW0056E: No default target service available for {0}**

**Explanation:** The default target service location cannot be obtained for the gateway service as no target services are defined.

**User Response:** Ensure that one or more target services are defined for the gateway service.

**WSGW0057E: No receiving channel name in context**

**Explanation:** A request has reached the gateway that does not contain the receiving channel name in the context.

**User Response:** Contact the supplier of the channel application.

**WSGW0058E: Channel {0} not defined for gateway service {1}**

**Explanation:** A request has reached the gateway for the given service through a channel which is not defined for that service. The request is rejected.

**User Response:** If the channel should be valid for the service, add the channel, otherwise check that the client of the request is making a valid request. This exception may be thrown when a client is making a malicious attack.

**WSGW0059E: gateway service {0} does not exist**

**Explanation:** A request was made for a gateway service that does not exist.

**User Response:** Ensure that the correct gateway service name is specified.

**WSGW0060E: gateway service {0} already exists**

**Explanation:** An attempt was made to create a new gateway service using a name that is used by an existing gateway service.

**User Response:** Use a different name for the gateway service.

**WSGW0061E: Could not find Service in UDDI registry {0} with parameters {1}, {2}, {3}**

**Explanation:** The given parameters for UDDI lookup did not yield a match.

**User Response:** Ensure that the parameters are correct. Also ensure that the UDDI reference parameters are correct and correspond to those used to publish the service to UDDI.

**WSGW0062E: Target service WSDL contains no <service> elements**

**Explanation:** The target service WSDL could be loaded but does not contain a <service> element. This is necessary to be able to invoke the target service.

**User Response:** Ensure that the target service WSDL contains one or more <service> element.

**WSGW0063E: Target service WSDL contains more than one service, and either target service name or namespace not specified**

**Explanation:** When adding a target service to a gateway service, you must specify both the service name and namespace values if there is more than one <service> element in the target service WSDL.

**User Response:** Specify the target service name and namespace as well as the location.

**WSGW0064E: Target service name {0} does not match service name in WSDL: {1}**

**Explanation:** A target service name was specified that is not the same as any target service name in the WSDL at the given location.

**User Response:** Ensure that a valid target service name is specified.

**WSGW0065E: Target service namespace {0} does not match service namespace in WSDL: {1}**

**Explanation:** A target service namespace was specified that is not the same as any target service namespace in the WSDL at the given location.

**User Response:** Ensure that a valid target service namespace is specified.

**WSGW0066E: Target service name {0} or namespace {1} not found in WSDL definition**

**Explanation:** A target service name and namespace were both specified, but do not match any target service name and namespace combination in the WSDL at the given location.

**User Response:** Ensure that a valid target service name and namespace combination is specified.

**WSGW0067E: UDDI reference {0} cannot be removed because it is being used by a deployed service**

**Explanation:** UDDI references can only be removed when they are not in use by gateway services.

**User Response:** Remove the UDDI reference from gateway services to which it is deployed before removing the UDDI reference.

**WSGW0068E: UDDI reference {0} already exists**

**Explanation:** The name specified for the UDDI reference is the same as that of a UDDI reference that is currently deployed.

**User Response:** Choose a different name for the UDDI reference, or remove the existing UDDI reference of the given name.

**WSGW0069E: UDDI reference {0} not found**

**Explanation:** No UDDI reference is currently deployed with the given name.

**User Response:** Use the name of a UDDI reference that is currently deployed.

**WSGW0070E: Invalid target service location type {0}**

**Explanation:** The location type for the target service is not a valid value.



- User Response:** Ensure that a correct value is specified for the target service location type.
- WSGW0071E: Failed to load URL definition from {0}**  
**Explanation:** The URL location specified was incorrect, or the WSDL it refers to cannot be loaded.
- User Response:** Ensure that the URL location is correct, and refers to a valid WSDL document.
- WSGW0072E: Failed to load UDDI definition from {0}**  
**Explanation:** The UDDI location specified was incorrect, or the WSDL it refers to cannot be loaded.
- User Response:** Ensure that the UDDI location is correct, and refers to a valid WSDL document.
- WSGW0073W: Not used**  
**Explanation:**
- User Response:**
- WSGW0074W: Invocation of filter {0} failed. Exception: {1}**  
**Explanation:** An unexpected exception was thrown during processing of the named filter.
- User Response:** Investigate the reason for the exception being thrown. Ensure that the filter application is available.
- WSGW0075E: Failed to set gateway end point address. Exception {0}**  
**Explanation:** An unexpected exception occurred when automatically setting the gateway's end point address.
- User Response:** Contact IBM Support
- WSGW0076E: Unable to access the gateway configuration bean. Exception {0}**  
**Explanation:** An unexpected exception occurred looking up the gateway's configuration bean in JNDI.
- User Response:** Restart the application server.
- WSGW0077E: Failed to remove gateway configuration session. Exception {0}**  
**Explanation:** An unexpected exception occurred removing the session bean while access the gateway's configuration bean.
- User Response:** Contact IBM Support
- WSGW0078E: Unable to access the gateway EndPoint bean. Exception {0}**  
**Explanation:** An unexpected exception occurred looking up the gateway's endpoint bean in JNDI.
- User Response:** Restart the application server.
- WSGW0079E: Failed to remove endpoint session. Exception {0}**  
**Explanation:** An unexpected exception occurred removing the session bean while access the gateway's endpoint bean.
- User Response:** Contact IBM Support
- WSGW0080E: Performance monitoring error. Exception {0}**  
**Explanation:** An unexpected exception occurred when recording performance monitoring information.
- User Response:** Contact IBM Support
- WSGW0081E: Unexpected error in method {0}. Exception {1}**  
**Explanation:** An unexpected exception occurred in the given method.
- User Response:** Contact IBM Support
- WSGW0082E: Unable to determine WAS security setting**  
**Explanation:** The WAS security setting could not be determined. It will be assumed that security is enabled.
- User Response:** No action required.
- WSGW0083W: Failed to authorize request for operation {0} on service {1}. Exception {2}**  
**Explanation:** Authorization of the given request failed. The request has been rejected.
- User Response:** Ensure that the required authorization bean has been generated for the given service, and that the correct authorization policy is defined.

**WSGW0084W: Invocation of filter {0} version {1} failed. Exception {2}**

**Explanation:** An exception was thrown during processing of the given filter. Processing of the request continues.

**User Response:** Investigate the reason for the exception being thrown by the filter. Refer to the documentation for the filter on how to resolve the problem.

**WSGW0085E: Failed to publish service {0} to UDDI registry {1}. Exception: {2}**

**Explanation:** An unexpected exception occurred when publishing the given service to a UDDI registry.

**User Response:** Ensure that the properties of the gateway service and UDDI reference are specified correctly.

**WSGW0086E: Failed to unpublish service {0} from UDDI registry {1}. Exception: {2}**

**Explanation:** An unexpected exception occurred when unpublishing the given service from a UDDI registry.

**User Response:** Ensure that the properties of the gateway service and UDDI reference are specified correctly.

**WSGW0087I: Published service {0} to UDDI registry {1}**

**Explanation:** The service was successfully published to the UDDI registry.

**User Response:** None

**WSGW0088I: Unpublished service {0} from UDDI registry {1}**

**Explanation:** The service was successfully unpublished from the UDDI registry.

**User Response:** None

**WSGW0089I: No MessageWarehouse registered. Requests will not be logged**

**Explanation:** A MessageWarehouse implementation was not found at the expected location in JNDI, so none is being used.

**User Response:** If a MessageWarehouse has been implemented, ensure that it is bound to JNDI at the correct location.

**WSGW0090I: No ExceptionHandler registered. Exceptions will not be handled**

**Explanation:** An ExceptionHandler implementation was not found at the expected location in JNDI, so none is being used.

**User Response:** If an ExceptionHandler has been implemented, ensure that it is bound to JNDI at the correct location.

**WSGW0091I: Usage: java -jar GenAuth -DWAS\_HOME=<was.install.directory> <HostName> <ServiceName>**

where <was.install.directory> is the location of the WebSphere installation directory and <HostName> is the url pointed to by the installation of the gateway and <ServiceName> is the name of the deployed gateway service. (Please note the ServiceName is case sensitive).

For example

```
-jar GenAuth.jar -DWAS_HOME=c:\\websphere\\AppServer http://host.machine.name.com/wsgw ServiceName
```

Successful execution will generate a file named <ServiceName>.ear

**Explanation:** Usage statement. This message is used by the WSGWAuthGen command-line utility.

**User Response:** No action required.

**WSGW0092I: Retrieving Service :**

**Explanation:** Progress message indicating that the service definition is being retrieved. This message is used by the WSGWAuthGen command line utility.

**User Response:** No action required.



**WSGW0093I: Retrieving Port Type :**

**Explanation:** Progress message indicating that the port type information is being retrieved. This message is used by the WSGWAuthGen command line utility.

**User Response:** No action required.

**WSGW0094I: Retrieving Methods :**

**Explanation:** Progress message indicating that method information is being retrieved. This message is used by the WSGWAuthGen command line utility.

**User Response:** No action required.

**WSGW0095I: Making Directory :**

**Explanation:** Progress message indicating that a directory is being created. This message is used by the WSGWAuthGen command-line utility.

**User Response:** No action required.

**WSGW0096I: Using Directory :**

**Explanation:** Progress message indicating that a directory is being used. This message is used by the WSGWAuthGen command-line utility.

**User Response:** No action required.

**WSGW0097I: About to compile....**

**Explanation:** Progress message indicating that a compilation is about to start. This message is used by the WSGWAuthGen command line utility.

**User Response:** No action required.

**WSGW0098I: Command Status :**

**Explanation:** General command status message. This message is used by the WSGWAuthGen command-line utility.

**User Response:** No action required.

**WSGW0099I: About to create jar....**

**Explanation:** Progress message indicating that a JAR file is about to be created. This message is used by the WSGWAuthGen command line utility.

**User Response:** No action required.

**WSGW0100I: About to create ear....**

**Explanation:** Progress message indicating that an EAR file is about to be created. This message is used by the WSGWAuthGen command line utility.

**User Response:** No action required.

**WSGW0101E: Error retrieving port from service {1}**

**Explanation:** An error occurred retrieving the port from the service in the WSDL. This message is used by the WSGWAuthGen command line utility.

**User Response:** Ensure that the service name is specified correctly and is deployed to the gateway with at least one target service and one channel.

**WSGW0102E: Error retrieving service {0}**

**Explanation:** An error occurred retrieving the service. This message is used by the WSGWAuthGen command-line utility.

**User Response:** Ensure that the service name is specified correctly and is deployed to the gateway with at least one target service and one channel.

**WSGW0103E: Exception while retrieving service definition from URL:**

**{0}/ServiceDefinition?name={1}. Exception: {2}**

**Explanation:** An unexpected exception occurred retrieving WSDL from the given location. This message is used by the WSGWAuthGen command-line utility.

**User Response:** Ensure that the service name is specified correctly and is deployed to the gateway with at least one target service and one channel.

**WSGW0104E: Error retrieving methods from service {0}**

**Explanation:** An unexpected exception occurred retrieving the methods that correspond to operations on the service.

**User Response:** Contact IBM Support

**WSGW0105E: Error retrieving WAS\_HOME environment variable**

**Explanation:** The value of the WAS\_HOME environment variable could not be retrieved.

**User Response:** Ensure that the WAS\_HOME variable is set correctly in the environment under which the command is being executed.

**WSGW0106E: Error compiling files**

**Explanation:** An unexpected error occurred compiling the generated Java files.

**User Response:** Contact IBM Support

**WSGW0107E: Error executing JAR command**

**Explanation:** An unexpected error occurred generating a JAR file.

**User Response:** Contact IBM Support

**WSGW0110E: A client attempted to load imported URL {0} for gateway service {1}. This URL is not imported by the definition for that service.**

**Explanation:** An attempt was made to use the gateway's import mapping servlet to load information from a URL that does not correspond to one that is referenced by the WSDL definition for that service.

**User Response:** Ensure that the client is making a valid request. This may be a malicious attempt to obtain information that the client does not have access to.

**WSGW0111W: Unsupported elements within the WSDL definition for target service {0} were ignored. The functionality of this service may be compromised.**

**Explanation:** In order to be able to use the given WSDL definition within the gateway, certain elements of the definition were ignored.

**User Response:** Refer to the service provider's documentation to determine whether this will affect the use of the service.

**WSGW0112E: Exception while sending response to queue {0}. Exception {1}, linked exception {2}**

**Explanation:** An unexpected exception occurred during execution of the JMS channel.

**User Response:** Ensure that the JMS channel and its prerequisites are correctly configured. Examine the linked exception for information on how to resolve the problem.

**WSGW0113E: Exception while creating Object Message to send response. Exception {1}, linked exception {2}**

**Explanation:** An unexpected exception occurred during execution of the JMS channel.

**User Response:** Ensure that the JMS channel and its prerequisites are correctly configured. Examine the linked exception for information on how to resolve the problem.

**WSGW0114E: Exception while creating Queue Sender to send response. Exception {1}, linked exception {2}**

**Explanation:** An unexpected exception occurred during execution of the JMS channel.

**User Response:** Ensure that the JMS channel and its prerequisites are correctly configured. Examine the linked exception for information on how to resolve the problem.

**WSGW0115E: Exception while creating Reply To Queue to send response. Exception {1}, linked exception {2}**

**Explanation:** An unexpected exception occurred during execution of the JMS channel.

**User Response:** Ensure that the JMS channel and its prerequisites are correctly configured. Examine the linked exception for information on how to resolve the problem.

**WSGW0116E: Exception while creating Queue Session to send response. Exception {1}, linked exception {2}**

**Explanation:** An unexpected exception occurred during execution of the JMS channel.

**User Response:** Ensure that the JMS channel and its prerequisites are correctly configured.  
Examine the linked exception for information on how to resolve the problem.

**WSGW0117E: Exception while creating Queue Connection Factory to send response. Exception {1}, linked exception {2}**

**Explanation:** An unexpected exception occurred during execution of the JMS channel.

**User Response:** Ensure that the JMS channel and its prerequisites are correctly configured.  
Examine the linked exception for information on how to resolve the problem.

**WSGW0118E: End point address keyword {0} is not valid for the JMS channel**

**Explanation:** The keyword specified in the end point address for the JMS channel is not valid.

**User Response:** Consult the documentation for a list of valid keywords for the JMS channel end point address.

**WSGW0119E: End point address for the JMS channel must be keyword=value pairs separated by the '&' character**

**Explanation:** The end point address specified does not have the correct syntax.

**User Response:** Consult the documentation for a specification of the syntax of the JMS channel end point address.

**WSGW0120E: Not used in this version of the gateway**

**Explanation:**

**User Response:**

**WSGW0121E: Not used in this version of the gateway**

**Explanation:**

**User Response:**

**WSGW0122E: Not used in this version of the gateway**

**Explanation:**

**User Response:**

**WSGW0123E: Not used in this version of the gateway**

**Explanation:**

**User Response:**

**WSGW0124E: Failed to find Business in UDDI Register**

**Explanation:** Used for SRPUDDIPublisher.

**User Response:**

**WSGW0125E: Not used**

**Explanation:**

**User Response:**

**WSGW0126E: Not used**

**Explanation:**

**User Response:**

**WSGW0127E: Not used**

**Explanation:**

**User Response:**

**WSGW0128E: Not used in this version of the gateway**

**Explanation:**

**User Response:**

**WSGW0129E: Not used in this version of the gateway**

**Explanation:**

**User Response:**

**WSGW0130E: Not used in this version of the gateway**

**Explanation:**

**User Response:**

**WSGW0131E: Not used in this version of the gateway**  
**Explanation:**  
**User Response:**

**WSGW0132E: Not used in this version of the gateway**  
**Explanation:**  
**User Response:**

**WSGW0133E: Not used in this version of the gateway**  
**Explanation:**  
**User Response:**

**WSGW0134E: Not used in this version of the gateway**  
**Explanation:**  
**User Response:**

**WSGW0135E: Not used in this version of the gateway**  
**Explanation:**  
**User Response:**

**WSGW0136E: Not used in this version of the gateway**  
**Explanation:**  
**User Response:**

**WSGW0137E: Not used in this version of the gateway**  
**Explanation:**  
**User Response:**

**WSGW0138E: Not used in this version of the gateway**  
**Explanation:**  
**User Response:**

**WSGW0139E: Not used in this version of the gateway**  
**Explanation:**  
**User Response:**

**WSGW0140E: Not used in this version of the gateway**  
**Explanation:**  
**User Response:**

**WSGW0141E: Not used in this version of the gateway**  
**Explanation:**  
**User Response:**

**WSGW0142E: Not used in this version of the gateway**  
**Explanation:**  
**User Response:**

**WSGW0143E: Not used in this version of the gateway**  
**Explanation:**  
**User Response:**

**WSGW0144E: Not used in this version of the gateway**  
**Explanation:**  
**User Response:**

**WSGW0145E: Not used in this version of the gateway**  
**Explanation:**  
**User Response:**

**WSGW0146E: Not used in this version of the gateway**  
**Explanation:**  
**User Response:**

**WSGW0147E: Not used in this version of the gateway**

**Explanation:**

**User Response:**

**WSGW0148E: A problem occurred whilst processing an attachment**

**Explanation:**

**User Response:**

**WSGW0149E: Failed to import the configuration. Exception: {0}**

**Explanation:** An unexpected error occurred when trying to import the configuration.

**User Response:** Contact IBM Support

**WSGW0150W: Security config for gateway service {0} has errors. Service will not be available.**

**Exception: {1}**

**Explanation:** The security configuration for the gateway service has errors which prevent the service from being available.

**User Response:** Review the security settings for the gateway service.

**WSGW0151W: Security config for target service {0} has errors. No security will be applied to the outgoing message. Exception: {1}**

**Explanation:** The security configuration for the target service has errors. No security will be applied to the outgoing message.

**User Response:** Review the security settings for the target service.

**WSGW0152E: A problem occurred whilst processing an attachment. Exception: {0}**

**Explanation:**

**User Response:**

**WSGW0153E: Handler name {0} already exists**

**Explanation:** The name specified for the handler is the same as that of a handler that is currently deployed.

**User Response:** Choose a different name for the handler, or remove the existing handler of the given name.

**WSGW0154E: Handler name {0} not found**

**Explanation:** No handler is currently deployed with the given name.

**User Response:** Use the name of a handler that is currently deployed.

**WSGW0155E: Error managing inbound handlers list at position {0} for service {2} and channel {1}**

**Explanation:**

**User Response:**

**WSGW0156E: Error managing outbound handlers list at position {0} for service {2} and target port {1}**

**Explanation:**

**User Response:**

**WSGW0157E: Target port {0} invalid or handlers list for the port does not exist**

**Explanation:**

**User Response:**

**WSGW0158E: Error removing handler {0} from outbound chain**

**Explanation:**

**User Response:**

**WSGW0159E: Channel name {0} invalid or handlers list for channel does not exist**

**Explanation:**

**User Response:**

**WSGW0160E: Attempt to remove handler {0} without "force" flag**

**Explanation:**

**User Response:**

**WSGW0161E: Unable to load handler implementation: {0}**

**Explanation:** The class specified could not be found on the classpath.

**User Response:** Make sure the class is available to the gateway application.

**WSGWJ001: An error occurred using the Web Services gateway.**

**Explanation:** This message is used by JSPs.

**User Response:**

**WSGWJ002: The following error occurred while processing the operation**

**Explanation:** This message is used by JSPs.

**User Response:**

**WSGWJ003: The reason this error occurred (the original exception) is as follows**

**Explanation:** This message is used by JSPs.

**User Response:**

**WSGWJ004: The stack trace for the error is as follows**

**Explanation:** This message is used by JSPs.

**User Response:**

**WSGWJ005: The stack trace for the original exception is as follows**

**Explanation:** This message is used by JSPs.

**User Response:**





---

## Chapter 13. Class loading

Class loaders affect the packaging of applications and the run-time behavior of packaged applications deployed on application servers.

1. Read about class loaders. The article "“Class loading: Resources for learning” on page 928” points to additional sources.
2. If necessary, migrate class-loader Module Visibility Mode settings for Version 4.0.x applications to Version 5.0 application or WAR class-loader policies.
3. If an application module uses a resource, create a resource provider that specifies the directory name of the resource drivers. Do not specify the resource JAR file names. All JAR files in the specified directory will be added into the class path of the WebSphere Application Server extensions class loader.

4. Configure class loaders of an application server for the run-time environment.

- a. Click **Servers > Application Servers > *server\_name*** and, on the settings page for an application server, set the application class-loader policy and application class-loader mode.

The application class-loader policy controls the isolation of applications running in the system. When set to SINGLE, applications are not isolated; a single application class loader is used to contain all EJB modules, dependency JAR files, and shared libraries in the system. When set to MULTIPLE, applications are isolated from each other; each application receives its own class loader to load that application’s EJB modules, dependency JAR files, and shared libraries.

The application class-loader mode specifies the class-loader mode when the application class-loader policy is SINGLE. PARENT\_FIRST causes the class loader to first delegate the loading of classes to its parent class loader before attempting to load the class from its local class path. PARENT\_LAST causes the class loader to first attempt to load classes from its local class path before delegating the class loading to its parent. This allows an application class loader to override and provide its own version of a class that exists in the parent class loader.

- b. On the settings page for an application server, click **Classloader**. On the Classloader page, click **New**.
  - c. On the settings page for a class loader, specify the class-loader mode. PARENT\_FIRST causes the class loader to delegate the loading of classes to its parent class loader before attempting to load the class from its local classpath. PARENT\_LAST causes the class loader to attempt to load classes from its local class path before delegating the class loading to its parent. Then, click **OK**.
  - d. On the settings page for a class loader, click **Libraries**. From the Library Ref page, click **Add**. On the settings page for a library reference, specify variables for the library reference as needed and click **OK**. Repeat the previous step until you define a library reference instance for each library file that your application needs. To define a library reference, you must first define one or more shared libraries. (See “Managing shared libraries” in the Information Center.)
5. When configuring an installed enterprise application for deployment in the run-time environment, set the class-loader mode and the WAR class-loader policy.
  6. When configuring an installed Web module for deployment in the run-time environment, set the class-loader mode.

---

### Class loaders

Class loaders are part of the Java virtual machine (JVM) code and are responsible for finding and loading class files. Class loaders affect the packaging of applications and the run-time behavior of packaged applications deployed on application servers.

The run-time environment of WebSphere Application Server uses the following class loaders to find and load new classes for an application in the following order:

1. The **bootstrap, extensions, and CLASSPATH class loaders** created by the JVM.

The bootstrap class loader uses the boot classpath (typically classes in jre/lib) to find and load classes. The extensions class loader uses the system property `java.ext.dirs` (typically `jre/lib/ext`) to find and load classes. The CLASSPATH class loader uses the CLASSPATH environment variable to find and load classes.

The CLASSPATH class loader contains the J2EE APIs of the WebSphere Application Server product (inside `j2ee.jar`). Because the J2EE APIs are in this class loader, you can add libraries that depend on J2EE APIs to the classpath system property to extend a server's classpath. However, a preferred method of extending a server's classpath is to add a shared library.

## 2. A **WebSphere-specific extensions class loader**.

The WebSphere extensions class loader loads the WebSphere run-time and J2EE classes that are required at run time. The extensions class loader uses a `ws.ext.dirs` system property to determine the path used to load classes. Each directory in the `ws.ext.dirs` classpath and every JAR file or ZIP file in these directories is added to the classpath used by this class loader.

The WebSphere extensions class loader also loads resource provider classes into a server if an application module installed on the server refers to a resource that is associated with the provider and if the provider specifies the directory name of the resource drivers.

## 3. One or more **application module class loaders** that load elements of enterprise applications running in the server.

The application elements can be Web modules, EJB modules, resource adapters, and dependency JAR files. Application class loaders follow J2EE class-loading rules to load classes and JAR files from an enterprise application. The WebSphere run time enables you to associate a shared library classpath with an application.

Each class loader is a child of the class loader above it. That is, the application module class loaders are children of the WebSphere-specific extensions class loader, which is a child of the CLASSPATH Java class loader. Whenever a class needs to be loaded, the class loader usually delegates the request to its parent class loader. If none of the parent class loaders can find the class, the original class loader attempts to load the class. Requests can only go to a parent class loader; they cannot go to a child class loader. If the WebSphere class loader is requested to find a class in a J2EE module, it cannot go to the application module class loader to find that class and a `ClassNotFoundException` occurs. Once a class is loaded by a class loader, any new classes that it tries to load reuse the same class loader or go up the precedence list until the class is found.

## **Class-loader isolation policies**

The number and function of the application module class loaders depends on the class-loader policies specified in the server configuration. Class loaders provide multiple options for isolating applications and modules to enable different application packaging schemes to run on an application server.

Two class-loader policies control the isolation of applications and modules:

### **Application class-loader policy**

Application class loaders consist of EJB modules, dependency JAR files, resource adapters, and shared libraries. Depending on the application class-loader policy, an application class loader can be shared by multiple applications (SINGLE) or unique for each application (MULTIPLE). The application class-loader policy controls the isolation of applications running in the system. When set to SINGLE, applications are not isolated. When set to MULTIPLE, applications are isolated from each other.

### **WAR class-loader policy**

By default, Web module class loaders load the contents of the `WEB-INF/classes` and `WEB-INF/lib` directories. The application class loader is the parent of the Web module class loader. You can change the default behavior by changing the application's WAR class-loader policy.

The WAR class-loader policy controls the isolation of Web modules. If this policy is set to APPLICATION, then the Web module contents also are loaded by the application class loader (in

addition to the EJB files, RAR files, dependency JAR files, and shared libraries). If the policy is set to MODULE, then each web module receives its own class loader whose parent is the application class loader.

**Note:** WebSphere server class loaders never load application client modules.


For each application server in the system, you can set the application class-loader policy to SINGLE or MULTIPLE. When the application class-loader policy is set to SINGLE, then a single application class loader loads all EJB modules, dependency JAR files, and shared libraries in the system. When the application class-loader policy is set to MULTIPLE, then each application receives its own class loader used for loading that application's EJB modules, dependency JAR files, and shared libraries.

This application class loader can load each application's Web modules if that WAR module's class-loader policy is also set to APPLICATION. If the WAR module's class-loader policy is set to APPLICATION, then the application's loader loads the WAR module's classes. If the WAR class-loader policy is set to MODULE, then each WAR module receives its own class loader.

The following example shows that when the application class-loader policy is set to SINGLE, a single application class loader loads all EJB modules, dependency JAR files, and shared libraries of all applications on the server. The single application class loader can also load Web modules if an application has its WAR class-loader policy set to APPLICATION. Applications having a WAR class-loader policy set to MODULE use a separate class loader for Web modules.

Application class-loader policy: SINGLE

```
Application 1
Module: EJB1.jar
Module: WAR1.war
 MANIFEST Class-Path: Dependency1.jar
 WAR Classloader Policy = MODULE
Application 2
Module: EJB2.jar
 MANIFEST Class-Path: Dependency2.jar
Module: WAR2.war
 WAR Classloader Policy = APPLICATION
```




The following example shows that when the application class-loader policy of an application server is set to MULTIPLE, each application on the server has its own class loader. An application class loader also loads its Web modules if the application's WAR class-loader policy is set to APPLICATION. If the policy is set to MODULE, then a Web module uses its own class loader.

Application class-loader policy: MULTIPLE

```

Application 1
Module: EJB1.jar
Module: WAR1.war
MANIFEST Class-Path: Dependency1.jar
WAR Classloader Policy = MODULE
Application 2
Module: EJB2.jar
MANIFEST Class-Path: Dependency2.jar
Module: WAR2.war
WAR Classloader Policy = APPLICATION

```



## Class-loader modes

There are two possible values for a class-loader mode:

### PARENT\_FIRST

The PARENT\_FIRST class-loader mode causes the class loader to first delegate the loading of classes to its parent class loader before attempting to load the class from its local classpath. This is the default for class-loader policy and for standard JVM class loaders.

### PARENT\_LAST

The PARENT\_LAST class-loader mode causes the class loader to first attempt to load classes from its local classpath before delegating the class loading to its parent. This policy allows an application class loader to override and provide its own version of a class that exists in the parent class loader.

The following settings determine a class loader's mode:

- If the application class-loader policy of an application server is SINGLE, the application class-loader policy of an application server defines the mode for an application class loader.
- If the application class-loader policy of an application server is MULTIPLE, the class-loader mode of an application defines the mode for an application class loader.

- If the WAR class-loader policy of an application is MODULE, the WAR class-loader policy of a Web module defines the mode for a WAR class loader.

---

## Class loader collection

Use this page to manage class-loader instances on an application server. A class loader determines whether an application class loader or a parent class loader finds and loads Java class files for an application.

To view this administrative console page, click **Servers > Application Servers > *server\_name* > Classloader**.

### Classloader ID

States a string unique to the server identifying the class-loader instance. The product assigns the identifier.

### Classloader Mode

Specifies the class-loader mode when the application class-loader policy is SINGLE. PARENT\_FIRST causes the class loader to delegate the loading of classes to its parent class loader before attempting to load the class from its local class path. PARENT\_LAST causes the class loader to attempt to load classes from its local class path before delegating the class loading to its parent; this allows an application class loader to override and provide its own version of a class that exists in the parent class loader.

## Class loader settings

Use this page to configure a class loader for applications that reside on an application server.

To view this administrative console page, click **Servers > Application Servers > *server\_name* > Classloader > *class\_loader\_ID***.

### Classloader ID

States a string unique to the server identifying the class-loader instance. The product assigns the identifier.

<b>Data type</b>	String
------------------	--------

### Classloader Mode

Specifies the class-loader mode when the application class-loader policy is SINGLE. PARENT\_FIRST causes the class loader to delegate the loading of classes to its parent class loader before attempting to load the class from its local class path. PARENT\_LAST causes the class loader to attempt to load classes from its local class path before delegating the class loading to its parent; this allows an application class loader to override and provide its own version of a class that exists in the parent class loader.

<b>Data type</b>	String
<b>Default</b>	PARENT_FIRST

---

## Migrating the class-loader Module Visibility Mode setting

WebSphere Application Server Version 4.0.x had a server-wide configuration setting called **Module Visibility Mode**. For Version 5.0, you use application or WAR class-loader policies instead of module visibility modes. The Version 5.0 policies provide additional flexibility because you can configure applications running in a server for an application class-loader policy of SINGLE or MULTIPLE and for a WAR class-loader policy of APPLICATION or MODULE.

To migrate module visibility modes in your Version 4.0.x applications to their equivalents in Version 5.0, change the settings for your Version 4.0.x applications and modules to the Version 5.0 values shown in the

table below.

Version 4.0.x module visibility mode	Version 5.0 application class-loader policy	Version 5.0 WAR class-loader policy
Server	SINGLE	APPLICATION
Compatibility	SINGLE	MODULE
Application	MULTIPLE	APPLICATION
Module*	MULTIPLE	MODULE
J2EE	MULTIPLE	MODULE

\*There is no exact equivalent for the Version 4.0.x Module mode because it isolated EJB modules within an application.

---

## Class loading: Resources for learning

Use the following links to find relevant supplemental information about class loaders. The information resides on IBM and non-IBM Internet sites, whose sponsors control the technical accuracy of the information.

These links are provided for convenience. Often, the information is not specific to the IBM WebSphere Application Server product, but is useful all or in part for understanding the product. When possible, links are provided to technical papers and Redbooks that supplement the broad coverage of the release documentation with in-depth examinations of particular product areas.

For current information available from IBM Support on known problems and their resolution, see the IBM Support page.

IBM Support has documents that can save you time gathering information needed to resolve this problem. Before opening a PMR, see the IBM Support page.

View links to additional information about:

- Programming model and decisions
- Programming instructions and examples
- Programming specifications

### Programming model and decisions

- J2EE Class Loading Demystified
- Understanding J2EE Application Server Class Loading Architectures

### Programming instructions and examples

- Developing and Deploying Modular J2EE Applications with WebSphere Studio Application Developer and WebSphere Application Server
- IBM WebSphere Application Server Programming

### Programming specifications

- Sun's J2EE™ Platform Specification
- Sun's J2EE™ Extension Mechanism Architecture

---

## Chapter 14. Using EJB query

The EJB query language is used to specify a query over container-managed entity beans. The language is similar to SQL. An EJB query is independent of the bean's mapping to a persistent store.

An EJB query can be used in three situations:

- To define a finder method of an EJB entity bean.
- To define a select method of an EJB entity bean.
- To dynamically specify a query using the `executeQuery()` dynamic API.

Finder and select queries are specified in the bean's deployment descriptor using the `<ejb-ql>` tag. Queries specified in the deployment descriptor are compiled into SQL during deployment. Dynamic queries require the interface provided by WebSphere Application Server Enterprise.

WebSphere's EJB query language is compliant with the EJB QL defined in Sun's EJB 2.0 specification and has additional capabilities as listed in the topic [Comparison of EJB 2.0 specification and WebSphere Query Language](#).

In your WebSphere application, you can define an EJB query in the following ways:

- **Assembly Toolkit.** When defining an EJB 2.0 entity bean in an EJB deployment descriptor editor, on the **Beans** page click **Add** under **Queries** and, in the Add Finder Descriptor wizard, define a `find` or `ejbSelect` method. See the online [Assembly Toolkit information](#) for documentation on wizard options.
- **WebSphere Studio Application Developer.** When defining an entity bean, specify the `<ejb-ql>` tag for the finder or select method.
- **Dynamic query service.** Add the `executeQuery()` method to your application. The dynamic query API is provided as an Enterprise Extension to WebSphere Application Server.

Before using EJB query, familiarize yourself with query language concepts, starting with the topic, [EJB Query Language](#).

See the topic [Example: EJB queries](#).

---

### EJB query language

An EJB query is a string that contains the following elements:

- a `SELECT` clause that specifies the EJBs or values to return;
- a `FROM` clause that names the bean collections;
- an optional `WHERE` clause that contains search predicates over the collections;
- an optional `GROUP BY` and `HAVING` clause (see [Aggregation functions](#));
- an optional `ORDER BY` clause that specifies the ordering of the result collection.

The `SELECT` clause is optional in order to maintain compatibility with WebSphere Application Server Version 4.

Collections of entity beans are identified in EJB queries through the use of their abstract schema name in the query `FROM` clause.

The elements of EJB query language are discussed in more detail in the following related topics.



## Example: EJB queries

Here is an example EJB schema, followed by a set of example queries:

Table 9. DeptBean schema

Entity bean name (EJB name)	DeptEJB (not used in query)
Abstract schema name	DeptBean
Implementation class	com.acme.hr.deptBean (not used in query)
Persistent attributes (cmp fields)	<ul style="list-style-type: none"><li>deptno - Integer (key)</li><li>name - String</li><li>budget - BigDecimal</li></ul>
Relationships	<ul style="list-style-type: none"><li>emps - 1:Many with EmpEJB</li><li>mgr - Many:1 with EmpEJB</li></ul>

Table 10. EmpBean schema

Entity bean name (EJB name)	EmpEJB (not used in query)
Abstract schema name	EmpBean
Implementation class	com.acme.hr.empBean (not used in query)
Persistent attributes (cmp fields)	<ul style="list-style-type: none"><li>empid - Integer (key)</li><li>name - String</li><li>salary - BigDecimal</li><li>bonus - BigDecimal</li><li>hireDate - java.sql.Date</li><li>birthDate - java.util.Calendar</li><li>address - com.acme.hr.Address</li></ul>
Relationships	<ul style="list-style-type: none"><li>dept - Many:1 with DeptEJB</li><li>manages - 1:Many with DeptEJB</li></ul>

Address is a serializable object used as cmp field in EmpBean. The definition of address is as follows:

```
public class com.acme.hr.Address extends Object implements Serializable {
public String street;
public String state;
public String city;
public Integer zip;
public double distance (String start_location) { ... } ;
public String format () { ... } ;
}
```

The following query returns all departments:

```
SELECT OBJECT(d) FROM DeptBean d
```

The following query returns departments whose name begins with the letters "Web". Sort the result by name:

```
SELECT OBJECT(d) FROM DeptBean d WHERE d.name LIKE 'Web%' ORDER BY d.name
```

The keywords SELECT and FROM are shown in uppercase in the examples but are case insensitive. If a name used in a query is a reserved word, the name must be enclosed in double quotes to be used in the query. There is a list of reserved words later in this document. Identifiers enclosed in double quotes are case sensitive. This example shows how to use a cmp field that is a reserved word:

```
SELECT OBJECT(d) FROM DeptBean d WHERE d."select" > 5
```

The following query returns all employees who are managed by Bob. This example shows how to navigate relationships using a path expression:

```
SELECT OBJECT (e) FROM EmpBean e WHERE e.dept.mgr.name='Bob'
```

A query can contain a parameter which refers to the corresponding value of the finder or select method. Query parameters are numbered starting with 1:

```
SELECT OBJECT (e) FROM EmpBean e WHERE e.dept.mgr.name= ?1
```

This query shows navigation of a multivalued relationship and returns all departments that have an employee that earns at least 50000 but not more than 90000:

```
SELECT OBJECT(d) FROM DeptBean d, IN (d.emps) AS e
WHERE e.salary BETWEEN 50000 and 90000
```

There is a join operation implied in this query between each department object and its related collection of employees. If a department has no employees, the department does not appear in the result. If a department has more than one employee that earns more than 50000, that department appears multiple times in the result.

The following query eliminates the duplicate departments:

```
SELECT DISTINCT OBJECT(d) from DeptBean d, IN (d.emps) AS e WHERE e.salary > 50000
```

Find employees whose bonus is more than 40% of their salary:

```
SELECT OBJECT(e) FROM EmpBean e where e.bonus > 0.40 * e.salary
```

Find departments where the sum of salary and bonus of employees in the department exceeds the department budget:

```
SELECT OBJECT(d) FROM DeptBean d where d.budget <
(SELECT SUM(e.salary+e.bonus) FROM IN(d.emps) AS e)
```

A query can contain DB2 style date-time arithmetic expressions if you use java.sql.\* datatypes as CMP fields and your datastore is DB2. Find all employees who have worked at least 20 years as of January 1st, 2000:

```
SELECT OBJECT(e) FROM EmpBean e where year('2000-01-01' - e.hireDate) >= 20
```

If the datastore is not DB2 or if you prefer to use java.util.Calendar as the CMP field, then you can use the java millisecond value in queries. The following query finds all employees born before Jan 1, 1990:

```
SELECT OBJECT(e) FROM EmpBean e WHERE e.birthDate < 631180800232
```

Find departments with no employees:

```
SELECT OBJECT(d) from DeptBean d where d.emps IS EMPTY
```

Find all employees whose earn more than Bob:

```
SELECT OBJECT(e) FROM EmpBean e, EmpBean b
WHERE b.name = 'Bob' AND e.salary + e.bonus > b.salary + b.bonus
```

Find the employee with the largest bonus:

```
SELECT OBJECT(e) from EmpBean e WHERE e.bonus =
(SELECT MAX(e1.bonus) from EmpBean e1)
```

The above queries all return EJB objects. A finder method query must always return an EJB Object for the home. A select method query can in addition return CMP fields or other EJB Objects not belonging to the home.

The following would be valid select method queries for EmpBean. Return the manager for each department:

```
SELECT d.mgr FROM DeptBean d
```

Return department 42 manager's name:

```
SELECT d.mgr.name FROM DeptBean d WHERE d.deptno = 42
```

Return the names of employees in department 42:

```
SELECT e.name FROM EmpBean e WHERE e.dept.deptno=42
```

Another way to write the same query is:

```
SELECT e.name from DeptBean d, IN (d.emps) AS e WHERE d.deptno=42
```

Finder and select queries allow only a single CMP field or EJBObject in the SELECT clause.

The dynamic query api allows multiple expressions in the SELECT clause. The following query would be a valid dynamic query, but not a valid select or finder query:

```
SELECT e.name, e.salary+e.bonus as total_pay , object(e), e.dept.mgr
FROM EmpBean e
ORDER BY 2
```

The following dynamic query returns the number of employees in each department:

```
SELECT e.dept.deptno as department_number , count(*) as employee_count
FROM EmpBean e
GROUP BY by e.dept.deptno
ORDER BY 1
```

The dynamic query api allows queries that contain bean or value object methods:

```
SELECT object(e), e.address.format()
FROM EmpBean e EmpBean e
```

## FROM clause

The FROM clause specifies the collections of objects to which the query is to be applied. Each collection is identified either by an abstract schema name and an identification variable, called a range variable, or by a collection member declaration that identifies a multivalued relationship and an identification variable.

Conceptually, the semantics of the query is to first form a temporary collection of tuples R. Tuples are composed of elements from the collections identified in the FROM clause. Each tuple contains one element from each of the collections in the FROM clause. All possible combinations are formed subject to the constraints imposed by the collection member declarations. If any schema name identifies a collection for which there are no records in the persistent store, then the temporary collection R will be empty.

### Example: FROM clause

DeptBean contains records 10, 20 and 30 in the persistent store. EmpBean contains records 1, 2 and 3 that are related to department 10 and records 4, 5 that are related to department 20. Department 30 has no related employees.

```
FROM DeptBean d, EmpBean e
```

This forms a temporary collection R that contains 15 tuples.

```
FROM DeptBean d, DeptBean d1
```

This forms a temporary collection R that contains 9 tuples.

```
FROM DeptBean d, IN (d.emps) AS e
```

This forms a temporary collection R that contains 5 tuples. Department 30 because it contains no employees will not be in R. Department 10 will be contained in R three times and department 20 will be contained in R twice.

After forming the temporary collection the search conditions of the WHERE clause will be applied to R and this will yield a new temporary collection R1. The ORDER BY and SELECT clauses are applied to R1 to yield the final result set.

An identification variable is a variable declared in the FROM clause using the operator IN or the optional AS.

```
FROM DeptBean AS d, IN (d.emps) AS e
```

is equivalent to:

```
FROM DeptBean d, IN (d.emps) e
```

An identification variable that is declared to be an abstract schema name is called a range variable. In the query above "d" is a range variable. An identification variable that is declared to be a multivalued path expression is called a collection member declaration. "d" and "e" in the example above are collection member declarations.

Note that the following path expression is illegal as a collection member declaration because it is not multivalued:

```
e.dept.mgr
```

## Inheritance in EJB query

If an EJB inheritance hierarchy has been defined for an abstract schema, using the abstract schema name in a query statement implies the collection of objects for that abstract schema as well as all subtypes.

### Example: Inheritance

Suppose that bean ManagerBean is defined as a subtype of EmpBean and ExecutiveBean is a subtype of ManagerBean in an EJB inheritance hierarchy. The following query returns employees as well as managers and executives:

```
SELECT OBJECT(e) FROM EmpBean e
```

## Path expressions

An identification variable followed by the navigation operator ( . ) and a cmp or relationship name is a path expression.

A path expression that leads to a cmr field can be further navigated if the cmr field is single-valued. If the path expression leads to a multi-valued relationship, then the path expression is terminal and cannot be further navigated. If the path expression leads to a cmp field whose type is a value object, it is possible to navigate to attributes of the value object.

### Example: Value object

Assume that address is a cmp field for EmpBean, which is a value object.

```
SELECT object(e) FROM EmpBean e
WHERE e.address.distance('San Jose') < 10 and e.address.zip = 95037
```

It is best to use the composer pattern to map value object attributes to relational columns if you intend to search on value attributes. If you store value objects in serialized format, then each value object must be retrieved from the database and deserialized. Value object methods can only be done in dynamic queries.

A path expression can also navigate to a bean method. The method must be defined on either the remote or local bean interface. Methods can only be used in dynamic queries. You cannot mix both remote and local methods in a single query statement.

If the query contains remote methods, the dynamic query must be executed using the query remote interface. Using the query remote interface causes the query service to activate beans and create instances of the remote bean interface

Likewise, a query statement with local bean methods must be executed with the query local interface. This causes the query service to activate beans and local interface instances.

Do not use get methods to access cmp and cmr fields of a bean.

If a method has overloaded definitions, the overloaded methods must have different number of parameters.

Methods must have non-void return types and method arguments and return types must be either primitive types byte, short, int, long, float, double, boolean, char or wrapper types from the following list:

Byte, Short, Integer, Long, Float, Double, BigDecimal, String, Boolean, Character, java.util.Calendar, java.sql.Date, java.sql.Time, java.sql.Timestamp, java.util.Date

If any input argument to a method is NULL, it is assumed the method returns a NULL value and the method is not invoked.

A collection valued path expression can be used in the FROM clause as a collection member declaration, and with the IS EMPTY, MEMBER OF, and EXISTS predicates in the WHERE clause.

FROM EmpBean e WHERE e.dept.mgr.name='Bob'	OK
FROM EmpBean e WHERE e.dept.emps.name='BOB'	INVALID -- cannot navigate through emps because it is multivalued
FROM EmpBean e, IN (e.dept.emps) e1 WHERE e1.name='BOB'	OK
FROM EmpBean e WHERE e.dept.emps IS EMPTY	OK

## WHERE clause

The WHERE clause contains search conditions composed of the following:

- literal values
- input parameters
- expressions
- basic predicates
- quantified predicates
- BETWEEN predicate
- IN predicate
- LIKE predicate
- NULL predicate
- EMPTY collection predicate
- MEMBER OF predicate
- EXISTS predicate
- IS OF TYPE predicate

If the search condition evaluates to TRUE, the tuple is added to the result set.

## Literals

A string literal is enclosed in single quotes. A single quote that occurs within a string literal is represented by two single quotes; For example: 'Tom's'. A string literal cannot exceed the maximum length that is supported by the underlying persistent datastore.

A numeric literal can be any of the following:

- an exact value such as 57, -957, +66
- any value supported by Java long
- a decimal literal such as 57.5, -47.02
- an approximate numeric value such as 7E3, -57.4E-2

A decimal or approximate numeric value must be in the range supported by the underlying persistent datastore.

A boolean literal can be the keyword TRUE or FALSE and is case insensitive.

## Input parameters

Input parameters are designated by the question mark followed by a number; For example: ?2

Input parameters are numbered starting at 1 and correspond to the arguments of the finder or select method; therefore, a query must not contain an input parameter that exceeds the number of input arguments.

An input parameter can be a primitive type of byte, short, int, long, float, double, boolean, char or wrapper types of Byte, Short, Integer, Long, Float, Double, BigDecimal, String, Boolean, Char, java.util.Calendar, java.util.Date, java.sql.Date, java.sql.Time, java.sql.Timestamp or an EJBObject.

An input parameter must not have a NULL value. To search for the occurrence of a NULL value the NULL predicate should be used.

## Expressions

Conditional expressions can consist of comparison operators and logical operators (AND, OR, NOT).

Arithmetic expressions can be used in comparison expressions and can be composed of arithmetic operations and functions, path expressions that evaluate to a numeric value and numeric literals and numeric input parameters.

String expressions can be used in comparison expressions and can be composed of string functions, path expressions that evaluate to a string value and string literals and string input parameters. A cmp field of type char is handled as if it were a string of length 1.

Boolean expressions can be used with = and <> comparison and can be composed of path expressions that evaluate to a boolean value and TRUE and FALSE keywords and boolean input parameters.

Reference expressions can be used with = and <> comparison and can be composed of path expressions that evaluate to a cmr field, an identification variable and an input parameter whose type is an EJB reference

Four different expression types are supported for working with date-time types. For portability the java.util.Calendar type should be used. DB2 style date, time and timestamp expressions are supported if the datastore is DB2 and the CMP field is of type java.util.Date, java.sql.Date, java.sql.Time or java.sql.Timestamp.

A Calendar type can be compared to another Calendar type, an exact numeric literal or input parameter of type long whose value is the standard Java long millisecond value.

The following query finds all employees born before Jan 1, 1990:

```
SELECT OBJECT(e) FROM EmpBean e WHERE e.birthDate < 631180800232
```

Date expressions can be used in comparison expressions and can be composed of operators + - , date duration expressions and date functions, path expressions that evaluate to a date value, string representation of a date and date input parameters.

Time expressions can be used in comparison expressions and can be composed of operators + - , time duration expressions and time functions, path expressions that evaluate to a time value, string representation of time and time input parameters.

Timestamp expressions can be used in comparison expressions and can be composed of operators + - , timestamp duration expressions and timestamp functions, path expressions that evaluate to a timestamp value, string representation of a timestamp and timestamp input parameters.

Standard bracketing ( ) for ordering expression evaluation is supported.

The operators and their precedence order from highest to lowest are:

- Navigation operator ( . )
- Arithmetic operators in precedence order:
  - + - unary
  - \* / multiply, divide
  - + - add, subtract
- Comparison operators: =, >, <, >=, <=, <>(not equal)
- Logical operator NOT
- Logical operator AND
- Logical operator OR

In some datastores, a zero length string value ( " ) is treated as a null value and affects the results of queries. Some datastores perform division between two integer values using integer arithmetic rules and other datastores use non integer rules. This also can affect the results of queries. For portability, avoid the use of zero length string values and division of integer values in an EJB query.

**Null value semantics:** The following describe the semantics of NULL values:

- Comparison or arithmetic operations with an unknown (NULL) value yield an unknown value
- Path expressions that contain NULL evaluate to NULL
- The IS NULL and IS NOT NULL operators can be applied to path expressions and return TRUE or FALSE. Boolean operators AND, OR and NOT use three valued logic.

AND	True	False	Unknown
True	True	False	Unknown
False	False	False	False
Unknown	Unknown	False	Unknown

OR	True	False	Unknown
True	True	True	True
False	True	False	Unknown
Unknown	True	Unknown	Unknown

	NOT
True	False
False	True
Unknown	Unknown



### Example: Null value semantics

```
select object(e) from EmpBean where e.salary > 10 and e.dept.budget > 100
```

If salary is NULL the evaluation of `e.salary > 10` returns unknown and the employee object is not returned. If the cmr field dept or budget is NULL evaluation of `e.dept.budget > 100` returns unknown and the employee object is not returned.

```
select object(e) from EmpBean where e.dept.budget is null
```

If dept or budget is NULL evaluation of `e.dept.budget is null` returns TRUE and the employee object is returned.

```
select object(e) from EmpBean e , in (e.dept.emps) e1 where e1.salary > 10
```

If dept is NULL, then the multivalued path expression `e.dept.emps` results in an empty collection (not a collection that contains a NULL value). An employee with a null dept value will not be returned.

```
select object(e) from EmpBean e where e.dept.emps is empty
```

If dept is NULL the evaluation of the predicate in unknown and the employee object is not returned.

```
select object(e) from EmpBean e , EmpBean e1 where e member of e1.dept.emps
```

If dept is NULL evaluation of the member of predicate returns unknown and the employee is not returned.

**Date time arithmetic and comparisons:** DATE, TIME and TIMESTAMP values may be compared with another value of the same type. Comparisons are chronological. Date time values can also be incremented, decremented, and subtracted.

If the datastore is DB2, then DB2 string representation of DATE, TIME and TIMESTAMP types can also be used. A string representation of a date or time can use ISO, USA, EUR or JIS format. A string representation of a timestamp uses ISO format.

Format	Date format	Date examples	Time format	Time examples
ISO	yyyy-mm-dd	1987-02-24 1987-2-24	hh.mm.ss	13.50.00 13.50
USA	mm/dd/yyyy	2/24/1987	hh:mm AM or PM	1:50 pm 02:10 AM
EUR	dd.mm.yyyy	24.02.1987 24.2.1987	hh.mm.ss	13.50.00 13.55
JIS	yyyy-mm-dd	1987-02-24	hh:mm:ss	13:50 13:50:05

### Example 1: Date time arithmetic comparisons

```
e.hiredate > '1990-02-24'
```

The timestamp of February 24th, 1990 1:50 pm can be represented as follows:

```
'1990-02-24-13.50.00.000000' or
'1990-02-24-13.50.00'
```

If the datastore is DB2, DB2 decimal durations can be used in expressions and comparisons. A date duration is a decimal(8,0) number that represents the difference between two dates in the format YYYYMMDD. A time duration is a decimal(6,0) number that represents the difference between two time values as HHMMSS. A timestamp duration is a decimal(20,6) number representing the differences between two timestamp values as YYYYMMDDHHMMSS.ZZZZZZ (ZZZZZZ is the number of microseconds and is to the right of the decimal point) .

Two date values (or time values or timestamp values) can be subtracted to yield a duration. If the second operand is greater than the first the duration is a negative decimal number. A duration can be added or subtracted from a datetime value to yield a new datetime value.

## Example 2: Date time arithmetic comparisons

`DATE('3/15/2000') - '12/31/1999'` results in a decimal number 215 which is a duration of 0 years, 2 months and 15 days.

Durations are really decimal numbers and can be used in arithmetic expressions and comparisons.

`( DATE('3/15/2000') - '12/31/1999' ) + 14 > 215` evaluates to TRUE.

`DATE('12/31/1999') + DECIMAL(215,8,0)` results in a date value 3/15/2000.

`TIME('11:02:26') - '00:32:56'` results in a decimal number 102930 which is a time duration of 10 hours, 29 minutes and 30 seconds.

`TIME('00:32:56') + DECIMAL(102930,6,0)` results in a time value of 11:02:26.

`TIME('00:00:59') + DECIMAL(240000,6,0)` results in a time value of 00:00:59.

`e.hiredate + DECIMAL(500,8,0) > '2000-10-01'` means compare the hiredate plus 5 months to the date 10/01/2000.

## Basic predicates

Basic predicates can be of two forms

`expression-1 comparison-operator expression-2`

`expression-3 comparison-operator ( subselect )`

The subselect must not return more than one value and the subselect can not return a type of an EJB reference. Boolean types and reference types only support = and <> comparisons.

### Example: Basic predicates

```
d.name='Java Development'
e.salary > 20000
e.salary > (select avg(e.salary) from EmpBean e)
```

## Quantified predicates

A quantified predicate compares a value with a set of values produced by a subselect.

`expression comparison-operator SOME | ANY | ALL ( subselect )`

The expression must not evaluate to a reference type.

When SOME or ANY is specified the result of the predicate is as follows:

- TRUE if the comparison is true for at least one value returned by the subselect.
- FALSE if the subselect is empty or if the comparison is false for every value returned by the subselect.
- UNKNOWN if the comparison is not true for all of the values returned by the subselect and at least one of the comparisons is unknown because of a null value.

When ALL is specified the result of the predicate is as follows:

- TRUE if the subselect returns empty or if the comparison is true to every value returned by the subselect.
- FALSE if the comparison is false for at least one value returned by the subselect.
- UNKNOWN if the comparison is not false for all values returned by the subselect and at least one comparison is unknown because of a null value.

## BETWEEN predicate

The BETWEEN predicate determines whether a given value lies between two other given values.

`expression [NOT] BETWEEN expression-2 AND expression-3`

**Example: BETWEEN predicate**

```
e.salary BETWEEN 50000 AND 60000
```

is equivalent to:

```
e.salary >= 50000 AND e.salary <= 60000
e.name NOT BETWEEN 'A' AND 'B'
```

is equivalent to:

```
e.name < 'A' OR e.name > 'B'
```

**IN predicate**

The IN predicate compares a value to a set of values and can have one of two forms:

```
expression [NOT] IN (subselect)
expression [NOT] IN (value1, value2,)
```

ValueN can either be a literal value or an input parameter. The expression can not evaluate to a reference type.

**Example: IN predicate**

```
e.salary IN (10000, 15000)
```

is equivalent to

```
(e.salary = 10000 OR e.salary = 15000)
e.salary IN (select e1.salary from EmpBean e1 where e1.dept.deptno = 10)
```

is equivalent to

```
e.salary = ANY (select e1.salary from EmpBean e1 where e1.dept.deptno = 10)
e.salary NOT IN (select e1.salary from EmpBean e1 where e1.dept.deptno = 10)
```

is equivalent to

```
e.salary <> ALL (select e1.salary from EmpBean e1 where e1.dept.deptno = 10)
```

**LIKE predicate**

The LIKE predicate searches a string value for a certain pattern.

```
string-expression [NOT] LIKE pattern [ESCAPE escape-character]
```

The pattern value is a string literal or parameter marker of type string in which the underscore ( `_` ) stands for any single character and percent ( `%` ) stands for any sequence of characters ( including empty sequence ). Any other character stands for itself. The escape character can be used to search for character `_` and `%`. The escape character can be specified as a string literal or an input parameter.

If the string-expression is null, then the result is unknown.

If both string-expression and pattern are empty, then the result is true.

**Example: LIKE predicate**

- `'' LIKE ''` is true
- `'' LIKE '%'` is true
- `e.name LIKE '12%3'` is true for '123' '12993' and false for '1234'
- `e.name LIKE 's_me'` is true for 'some' and 'same', false for 'soome'
- `e.name LIKE '/_foo'` escape '/' is true for '/\_foo', false for 'afoo'
- `e.name LIKE '//_foo'` escape '/' is true for '/afoo' and for '/bfoo'
- `e.name LIKE '///_foo'` escape '/' is true for '/\_foo' but false for '/afoo'

## NULL predicate

The NULL predicate tests for null values.

```
single-valued-path-expression IS [NOT] NULL
```

### Example: NULL predicate

```
e.name IS NULL
e.dept.name IS NOT NULL
e.dept IS NOT NULL
```

## EMPTY collection predicate

To test if a multivalued relationship is empty, use the following syntax:

```
collection-valued-path-expression IS [NOT] EMPTY
```

### Example: Empty collection predicate

To find all departments with no employees:

```
SELECT OBJECT(d) FROM DeptBean d WHERE d.emps IS EMPTY
```

## MEMBER OF predicate

This expression tests whether the object reference specified by the single valued path expression or input parameter is a member of the designated collection. If the collection valued path expression designates an empty collection the value of the MEMBER OF expression is FALSE.

```
{ single-valued-path-expression | input_parameter } [NOT] MEMBER [OF] collection-valued-path-expression
```

### Example: MEMBER OF predicate

Find employees that are not members of a given department number:

```
SELECT OBJECT(e) FROM EmpBean e , DeptBean d
WHERE e NOT MEMBER OF d.emps AND d.deptno = ?1
```

Find employees whose manager is a member of a given department number:

```
SELECT OBJECT(e) FROM EmpBean e, DeptBean d
WHERE e.dept.mgr MEMBER OF d.emps and d.deptno=?1
```

## EXISTS predicate

The exists predicate tests for the presence or absence of a condition specified by a subselect.

```
EXISTS (subselect)
EXISTS collection-valued-path-expression
```

The result of EXISTS is true if the subselect returns at least one value or the path expression evaluates to a nonempty collection, otherwise the result is false.

To negate an EXISTS predicate, precede it with the logical operator NOT.

### Example: EXISTS predicate

Return departments that have at least one employee earning more than 1000000:

```
SELECT OBJECT(d) FROM DeptBean d
WHERE EXISTS (SELECT 1 FROM IN (d.emps) e WHERE e.salary > 1000000)
```

Return departments that have no employees:

```
SELECT OBJECT(d) FROM DeptBean d
WHERE NOT EXISTS (SELECT 1 FROM IN (d.emps) e)
```

The above query can also be written as follows:

```
SELECT OBJECT(d) FROM DeptBean d WHERE NOT EXISTS d.emps
```

## IS OF TYPE predicate

The IS OF TYPE predicate is used to test the type of an EJB reference. It is similar in function to the Java instance of operator. IS OF TYPE is used when several abstract beans have been grouped into an EJB inheritance hierarchy. The type names specified in the predicate are the bean abstract names. The ONLY option can be used to specify that the reference must be exactly this type and not a subtype.

```
identification-variable IS OF TYPE ([ONLY] type-1, [ONLY] type-2,)
```

### Example: IS OF TYPE predicate

Suppose that bean ManagerBean is defined as a subtype of EmpBean and ExecutiveBean is a subtype of ManagerBean in an EJB inheritance hierarchy.

The following query returns employees as well as managers and executives:

```
SELECT OBJECT(e) FROM EmpBean e
```

If you are interested in objects which are employees and not managers and not executives:

```
SELECT OBJECT(e) FROM EmpBean e WHERE e IS OF TYPE(ONLY EmpBean)
```

If you are interested in object which are managers or executives:

```
SELECT OBJECT(e) FROM EmpBean e WHERE e IS OF TYPE(ManagerBean)
```

The above query is equivalent to the following query:

```
SELECT OBJECT(e) FROM ManagerBean e
```

If you are interested in managers only and not executives:

```
SELECT OBJECT(e) FROM EmpBean e WHERE e IS OF TYPE(ONLY ManagerBean)
```

or:

```
SELECT OBJECT(e) FROM ManagerBean e
WHERE e IS OF TYPE (ONLY ManagerBean)
```

## Scalar functions

EJB query contains scalar built-in functions for doing type conversions, string manipulation, and for manipulating date-time values. The list of scalar functions is documented in the topic EJB query: Scalar functions.

### Example: Scalar functions

Find employees hired in 1999:

```
SELECT OBJECT(e) FROM EmpBean e where YEAR(e.hireDate) = 1999
```

The only scalar functions that are guaranteed to be portable across backend datastore vendors are the following:

- ABS
- SQRT
- CONCAT
- LENGTH
- LOCATE
- SUBSTRING

The other scalar functions should be used only when DB2 is the backend datastore.

## EJB query: Scalar functions

EJB query contains scalar built-in functions, as listed below, for doing type conversions, string manipulation, and for manipulating date-time values.

### Numeric functions

ABS ( < any numeric datatype > ) -> < any numeric datatype >  
SQRT ( < any numeric datatype > ) -> Double

### Type conversion functions

CHAR ( < any numeric datatype > ) -> string  
CHAR ( < string > ) -> string  
CHAR ( < any datetime datatype > [, Keyword k ] ) -> string

Datetime datatype is converted to its string representation in a format specified by the keyword k. The valid keywords values are ISO, USA, EUR or JIS. If k is not specified the default is ISO.

BIGINT ( < any numeric datatype > ) -> Long  
BIGINT ( < string > ) -> Long

The following function converts the argument to an integer n by truncation and returns the date that is n-1 days after January 1, 0001:

DATE ( < date string > ) -> Date  
DATE ( < any numeric datatype> ) -> Date

The following function returns date portion of a timestamp:

DATE( timestamp ) -> Date  
DATE ( < timestamp-string > ) -> Date

The following function converts number to decimal with optional precision p and scale s.

DECIMAL ( < any numeric datatype > [, p [, s ] ] ) -> Decimal

The following function converts string to decimal with optional precision p and scale s.

DECIMAL ( < string > [, p [, s ] ] ) -> Decimal  
DOUBLE ( < any numeric datatype > ) -> Double  
DOUBLE ( < string > ) -> Double  
FLOAT ( < any numeric datatype > ) -> Double  
FLOAT ( < string > ) -> Double

Float is a synonym for DOUBLE.

INTEGER ( < any numeric datatype > ) -> Integer  
INTEGER ( < string > ) -> Integer  
REAL ( < any numeric datatype > ) -> Float  
SMALLINT ( < any numeric datatype > ) -> Short  
SMALLINT ( < string > ) -> Short  
TIME ( < time > ) -> Time  
TIME ( < time-string > ) -> Time  
TIME ( < timestamp > ) -> Time  
TIME ( < timestamp-string > ) -> Time  
TIMESTAMP ( < timestamp > ) -> Timestamp  
TIMESTAMP ( < timestamp-string > ) -> Timestamp

### String functions

CONCAT ( <string>, <string> ) -> String

The following function returns a character string representing absolute value of the argument not including its sign or decimal point. For example, digits( -42.35) is "4235".

DIGITS ( Decimal d ) -> String

The following function returns the length of the argument in bytes. If the argument is a numeric or datetime type, it returns the length of internal representation.

LENGTH ( < string > ) -> Integer

The following function returns a copy of the argument string where all upper case characters have been converted to lower case.

LCASE ( < string > ) -> String

The following function returns the starting position of the first occurrence of argument 1 inside argument 2 with optional start position. If not found, it returns 0.

LOCATE ( String s1 , String s2 [, Integer start ] ) -> Integer

The following function returns a substring of s beginning at character m and containing n characters. If n is omitted, the substring contains the remainder of string s. The result string is padded with blanks if needed to make a string of length n.

SUBSTRING ( String s , Integer m [ , Integer n ] ) -> String

The following function returns a copy of the argument string where all lower case characters have been converted to upper case.

UCASE ( < string > ) -> String

### **Date - time functions**

The following function returns the day portion of its argument. For a duration, the return value can be -99 to 99.

DAY ( Date ) -> Integer  
DAY ( < date-string > ) -> Integer  
DAY ( < date-duration > ) -> Integer  
DAY ( Timestamp ) -> Integer  
DAY ( < timestamp-string > ) -> Integer  
DAY ( < timestamp-duration > ) -> Integer

The following function returns one more than number of days from January 1, 0001 to its argument.

DAYS ( Date ) -> Integer  
DAYS ( < Date-string > ) -> Integer  
DAYS ( Timestamp ) -> Integer  
DAYS ( < timestamp-string > ) -> Integer

The following function returns the hour part of its argument. For a duration, the return value can be -99 to 99.

HOUR ( Time ) -> Integer  
HOUR ( < time-string > ) -> Integer  
HOUR ( < time-duration > ) -> Integer  
HOUR ( Timestamp ) -> Integer  
HOUR ( < timestamp-string > ) -> Integer  
HOUR ( < timestamp-duration > ) -> Integer

The following function returns the microsecond part of its argument.

MICROSECOND ( Timestamp ) -> Integer  
MICROSECOND ( < timestamp-string > ) -> Integer  
MICROSECOND ( < timestamp-duration > ) -> Integer

The following function returns the minute part of its argument. For a duration, the return value can be -99 to 99.



```

MINUTE (Time) -> Integer
MINUTE (< time-string >) -> Integer
MINUTE (< time-duration >) -> Integer
MINUTE (Timestamp) -> Integer
MINUTE (< timestamp-string >) -> Integer
MINUTE (< timestamp-duration >) -> Integer

```

The following function returns the month portion of its argument. For a duration, the return value can be -99 to 99.

```

MONTH (Date) -> Integer
MONTH (< date-string >) -> Integer
MONTH (< date-duration >) -> Integer
MONTH (Timestamp) -> Integer
MONTH (< timestamp-string >) -> Integer
MONTH (< timestamp-duration >) -> Integer

```

The following function returns the second part of its argument. For a duration, the return value can be -99 to 99.

```

SECOND (Time) -> Integer
SECOND (< time-string >) -> Integer
SECOND (< time-duration >) -> Integer
SECOND (Timestamp) -> Integer
SECOND (< timestamp-string >) -> Integer
SECOND (< timestamp-duration >) -> Integer

```

The following function returns the year portion of its argument. For a duration, the return value can be -9999 to 9999.

```

YEAR (Date) -> Integer
YEAR (< date-string >) -> Integer
YEAR (< date-duration >) -> Integer
YEAR (Timestamp) -> Integer
YEAR (< timestamp-string >) -> Integer
YEAR (< timestamp-duration >) -> Integer

```

## Aggregation functions

Queries that return aggregate values can only be used with the dynamic query interface available in WebSphere Business Integration Server. However, aggregation functions can be used in non-dynamic queries if the aggregation function is used in a subselect or HAVING clause.

Aggregation functions operate on a set of values to return a single scalar value. The following is an example of an aggregation:

```
SELECT SUM (e.salary + e.bonus) FROM EmpBean e WHERE e.dept.deptno =20
```

This computes the total salary and bonus for department 20.

The aggregation functions are avg, count, max, min and sum. The syntax of an aggregation function is as follows:

```
aggregation-function ([ALL | DISTINCT] expression)
```

or:

```
COUNT(*)
```

The DISTINCT option eliminates duplicate values before applying the function. ALL is the default and does not eliminate duplicates. Null values are ignored in computing the aggregate function except for COUNT(\*) which returns a count of all elements in the set.

MAX and MIN can apply to any numeric, string or datetime datatype and returns the same datatype. SUM and AVG take a numeric type as input. SUM and AVG return numeric type. The actual numeric type returned by SUM and AVG depends on the underlying datastore. COUNT can take any datatype and returns an integer.

If you are using an Informix datastore, the argument to COUNT must be an asterisk or a single valued path expression. The argument to SUM, AVG, MIN, or MAX used with DISTINCT must be a single valued path expression.

The set of values that is used for the aggregate function is determined by the collection that results from the FROM and WHERE clause of the subquery. This set can be divided into groups and the aggregation function applied to each group. This is done by using a GROUP BY clause in the subquery. The GROUP BY clause defines grouping members which is a list of path expressions. Each path expression specifies a field that is a primitive type of byte, short, int, long, float, double, boolean, char, or a wrapper type of Byte, Short, Integer, Long, Float, Double, BigDecimal, String, Boolean, Character, java.util.Calendar, java.util.Date, java.sql.Date, java.sql.Time or java.sql.Timestamp.

Finder or select queries can not return aggregation function values. In other words, aggregation functions can not appear in the top level SELECT of a finder or select query but can be used in subqueries.

### Example: Aggregation functions

```
SELECT e.dept.deptno, AVG (e.salary) FROM EmpBean e GROUP BY e.dept.deptno
```

The above query computes the average salary for each department.

In dividing a set into groups, a NULL value is considered equal to another NULL value.

Just as the WHERE clause filters tuples from the FROM clause, the groups can be filtered using a HAVING clause that tests group properties involving aggregate functions or grouping members:

```
SELECT e.dept.deptno, AVG (e.salary) FROM EmpBean e
GROUP BY e.dept.deptno
HAVING COUNT(*) > 3 AND e.dept.deptno > 5
```

This query returns average salary for departments that have more than 3 employees and the department number is greater than 5.

It is possible to have a HAVING clause without a GROUP BY clause in which case the entire set is treated as a single group to which the HAVING clause is applied.

## SELECT clause

For finder and select queries, the syntax of the SELECT clause is as follows:

```
SELECT [ALL | DISTINCT]
{ single-valued-path-expression | OBJECT (identification-variable) }
```

The SELECT clause consists of either a single identification variable that is defined in the FROM clause or a single valued path expression that evaluates to a object reference or CMP value. The keyword DISTINCT can be used to eliminate duplicate references.

For a query that defines a finder method the query must return an object type consistent with the home for which the finder method associated with the query. In other words, a finder method for a department home can not return employee objects.

For dynamic queries the syntax is as follows:

```
SELECT { ALL | DISTINCT } [selection ,]* selection
selection ::= { expression [[AS] id] | scalar-subselect }
```

A scalar-subselect is a subselect that returns a single value.

### Example: SELECT clause

Find all employees that earn more than John:

```
SELECT OBJECT(e) FROM EmpBean ej, EmpBean e
WHERE ej.name = 'John' and e.salary > ej.salary
```

Find all departments that have one or more employees who earn less than 20000:

```
SELECT DISTINCT e.dept FROM EmpBean e where e.salary < 20000
```

A select method query can have a path expression that evaluates to an arbitrary value:

```
SELECT e.dept.name FROM EmpBean e where e.salary < 2000
```

The above query returns a collection of name values for those departments having employees earning less than 20000.

### Example: Valid dynamic queries

The following are examples of dynamic queries:

```
SELECT e.name, e.salary+e.bonus as total_pay from EmpBean e
SELECT SUM(e.salary+e.bonus) from EmpBean e where e.dept.deptno = ?1
```

## ORDER BY clause

The ORDER BY clause specifies an ordering of the objects in the result collection:

```
ORDER BY [order_element ,]* order_element
order_element ::= { path-expression | integer } [ASC | DESC]
```

The path expression must specify a single valued field that is a primitive type of byte, short, int, long, float, double, char or a wrapper type of Byte, Short, Integer, Long, Float, Double, BigDecimal, String, Character, java.util.Calendar, java.util.Date, java.sql.Date, java.sql.Time, java.sql.Timestamp.

ASC specifies ascending order and is the default. DESC specifies descending order.

Integer refers to a selection expression in the SELECT clause.

### Example: ORDER BY clause

Return department objects in decreasing deptno order:

```
SELECT OBJECT(d) FROM DeptBean d ORDER BY d.deptno DESC
```

Return employee objects sorted by department number and name:

```
SELECT OBJECT(e) FROM EmpBean e ORDER BY e.dept.deptno ASC, e.name DESC
```

The following is a valid dynamic query:

```
SELECT OBJECT(e), e.salary+e.bonus as total_pay FROM EmpBean e ORDER BY 2 DESC
```

## Subqueries

A subquery can be used in quantified predicates, EXISTS predicate or IN predicate. A subquery should only specify a single element in the SELECT clause. When a path expression appears in a subquery, the identification variable of the path expression must be defined either in the subquery, in one of the containing subqueries, or in the outer query. A scalar subquery is a subquery that returns one value. A scalar subquery can be used in a basic predicate and in the SELECT clause of a dynamic query.

### Example: Subqueries

```
SELECT OBJECT(e) FROM EmpBean e
WHERE e.salary > (SELECT AVG(e1.salary) FROM EmpBean e1)
```

The above query returns employees who earn more than average salary of all employees.

```
SELECT OBJECT(e) FROM EmpBean e WHERE e.salary >
(SELECT AVG(e1.salary) FROM IN (e.dept.emps) e1)
```

The above query returns employees who earn more than average salary of their department.

```
SELECT OBJECT(e) FROM EmpBean e WHERE e.salary =
(SELECT MAX(e1.salary) FROM IN (e.dept.emps) e1)
```

The above query returns employees who earn the most in their department.

```
SELECT OBJECT(e) FROM EmpBean e
WHERE e.salary > (SELECT AVG(e.salary) FROM EmpBean e1
WHERE YEAR(e1.hireDate) = YEAR(e.hireDate))
```

The above query returns employees who earn more than the average of employees hired in same year.

## EJB query restrictions

An EJB query is compiled into an SQL query and executed against the underlying datastore based on schema mapping of the abstract bean to the datastore schema. The semantics of comparison and arithmetic operations are that of the underlying datastore. In the case of SQL, note that two strings are equal if the shorter string padded with blanks equals the longer string. For example, 'A' is equal to 'A '. This differs from the equality of strings in the Java language. Arithmetic overflow operations are an error in SQL.

A cmp field can not be used in comparison operations or predicates (except for LIKE) if that cmp field is mapped to a long varchar or large objects (LOB) column or any other column type for which the database server does not support predicates or comparison operations.

A cmp field of any type can be used in a SELECT clause. Fields that can be used in predicates, grouping, or ordering operations must be of the types listed below:

- Primitive types : byte, short, int, long, float, double, boolean, char
- Object types: Byte, Short, Integer, Long, Float, Double, BigDecimal, String, Boolean, Char, java.util.Calendar , java.util.Date
- JDBC types: java.sql.Date, java.sql.Time, java.sql.Timestamp

The field must be mapped to a table column that is compatible in type either by using a "top-down" default mapping generated by the WebSphere deploy tool, or using a "meet-in-the-middle" mapping between compatible types.

In order to search on attributes of a cmp field that is a user-defined value object, you should use a "meet-in-the-middle" mapping and use a composer to map each attribute to a compatible column. The default "top-down" mapping stores the object as a serialized object in a column of type blob, which does not allow searching.

If a cmp field is mapped to a column using a "meet-in-the-middle" mapping with a converter, that field can only be used with the NULL predicate or with basic predicates of the following form:

```
path-expression <comparison> literal_value
path-expression <comparison> input_parameter
```

In this situation, the converter method `toData()` is called to convert the right-hand side of the predicate to an SQL value.

Example of allowable predicate on a cmp field with user defined converter:

```
e.name = 'Chris'
e.name > ?1
e.name IS NULL
```

Examples of unallowable predicates:

```
substring(e.name, 1, 3) = 'ABC'
e.salary > d.budget
```

A converter should preserve equality, collating sequence and null values when doing a conversion. Otherwise cmp fields created by the converter should not be used in WHERE, GROUP, HAVING or ORDER clauses of a query.

## EJB Query: Reserved words

The following words are reserved in WebSphere EJB query:

all, as, distinct, empty, false, from, group, having, in, is, like, select, true, union, where

Avoid using identifiers that start with underscore (for example, `_integer`) as these are also reserved.

## EJB query: BNF syntax

```
EJB QL ::= [select_clause] from_clause [where_clause]
 [group_by_clause] [having_clause] [order_by_clause]
from_clause ::= FROM identification_variable_declaration
 [,identification_variable_declaration]*
identification_variable_declaration ::= collection_member_declaration |
 range_variable_declaration
collection_member_declaration ::=
 IN (collection_valued_path_expression) [AS] identifier
range_variable_declaration ::= abstract_schema_name [AS] identifier
single_valued_path_expression ::=
 {single_valued_navigation | identification_variable}. (cmp_field |
 method | cmp_field.value_object_attribute | cmp_field.value_object_method)
 | single_valued_navigation
single_valued_navigation ::=
 identification_variable.[single_valued_cmr_field.]*
 single_valued_cmr_field
collection_valued_path_expression ::=
 identification_variable.[single_valued_cmr_field.]*
 collection_valued_cmr_field
select_clause ::= SELECT { ALL | DISTINCT } {single_valued_path_expression |
 identification_variable | OBJECT (identification_variable) }
select_clause_eex ::= SELECT { ALL | DISTINCT } [selection ,]* selection
selection ::= { expression [[AS] id] | subselect }
order_by_clause ::= ORDER BY [{single_valued_path_expression | integer} [ASC|DESC],]*
 {single_valued_path_expression | integer}[ASC|DESC]
where_clause ::= WHERE conditional_expression
conditional_expression ::= conditional_term |
 conditional_expression OR conditional_term
conditional_term ::= conditional_factor |
 conditional_term AND conditional_factor
conditional_factor ::= [NOT] conditional_primary
conditional_primary ::= simple_cond_expression | (conditional_expression)
```

```

simple_cond_expression ::= comparison_expression | between_expression |
 like_expression | in_expression | null_comparison_expression |
 empty_collection_comparison_expression | quantified_expression |
 exists_expression | is_of_type_expression | collection_member_expression
between_expression ::= expression [NOT] BETWEEN expression AND expression
in_expression ::= single_valued_path_expression [NOT] IN
 { (subselect) | (atom ,)* atom }
atom = { string-literal | numeric-constant | input-parameter }
like_expression ::= expression [NOT] LIKE
 {string_literal | input_parameter}
 [ESCAPE {string_literal | input_parameter}]
null_comparison_expression ::=
 single_valued_path_expression IS [NOT] NULL
empty_collection_comparison_expression ::=
 collection_valued_path_expression IS [NOT] EMPTY
collection_member_expression ::=
 { single_valued_path_expression | input_paramter } [NOT] MEMBER [OF]
 collection_valued_path_expression
quantified_expression ::=
 expression comparison_operator {SOME | ANY | ALL} (subselect)
exists_expression ::= EXISTS {collection_valued_path_expression | (subselect)}
subselect ::= SELECT [{ ALL | DISTINCT }] expression from_clause [where_clause]
 [group_by_clause] [having_clause]
group_by_clause ::= GROUP BY [single_valued_path_expression,]*
 single_valued_path_expression
having_clause ::= HAVING conditional_expression
is_of_type_expression ::= identifier IS OF TYPE
 ([[ONLY] abstract_schema_name,]* [ONLY] abstract_schema_name)
comparison_expression ::= expression comparison_operator { expression | (subquery) }
comparison_operator ::= = | > | >= | < | <= | <>
method ::= method_name([[expression ,]* expression])
expression ::= term | expression {+|-} term
term ::= factor | term {*/|/} factor
factor ::= {+|-} primary
primary ::= single_valued_path_expression | literal |
 (expression) | input_parameter | functions
functions ::=
 ABS(expression) |
 AVG([ALL|DISTINCT] expression) |
 BIGINT(expression) |
 CHAR({expression [, {ISO|USA|EUR|JIS}] }) |
 CONCAT (expression , expression) |
 COUNT({[ALL|DISTINCT] expression | *}) |
 DATE(expression) |
 DAY({expression } |
 DAYS(expression) |
 DECIMAL(expression [,integer[,integer]])
 DIGITS(expression) |
 DOUBLE(expression) |
 FLOAT(expression) |
 HOUR (expression) |
 INTEGER(expression) |
 LCASE (expression) |
 LENGTH(expression) |
 LOCATE(expression, expression [, expression]) |
 MAX([ALL|DISTINCT] expression) |
 MICROSECOND(expression) |
 MIN([ALL|DISTINCT] expression) |
 MINUTE (expression) |

```

MONTH( expression ) |  
 REAL( expression ) |  
 SECOND( expression ) |  
 SMALLINT( expression ) |  
 SQRT ( expression ) |  
 SUBSTRING( expression, expression[, expression]) |  
 SUM([ALL|DISTINCT] expression) |  
 TIME( expression ) |  
 TIMESTAMP( expression ) |  
 UCASE ( expression ) |  
 YEAR( expression )

## Comparison of EJB 2.0 specification and WebSphere query language

Item	EJB 2.0 specification	WebSphere Query	WebSphere Enterprise (Dynamic) Query
Bean methods	no	no	yes
Calendar comparisons	yes	yes	yes
Delimited identifiers	no	yes	yes
Dependent Value attributes	no	yes	yes
Dependent Value methods	no	no	yes
Dynamic Query APIs	no	no	yes
EXISTS predicate	no	yes	yes
Inheritance	no	yes	yes
Multiple element select clauses	no	no	yes
Order by	no	yes	yes
Scalar functions	yes *	yes	yes
Select clause	required	optional	optional
SQL Date/time expressions	no	yes	yes
String comparisons	= and <> only	= <> > <	= <> > <
Subqueries, aggregations, group by, and having clauses	no	yes	yes

\* EJB 2.0 defines the following scalar functions: abs, sqrt, concat, length, locate and substring. WebSphere query and dynamic query support additional scalar functions as listed in the topic, EJB query: Scalar functions.

---

## Using the dynamic query service

Consider using the dynamic query service (available with Integration Server) when any of the following are true:

- You do not know the query search criteria until application runtime.
- You need to return multiple cmp or cmr fields from a query (deployment queries allow only a single element to be specified in the SELECT clause). For more information, see the Example: EJB queries article.
- You want to perform aggregation in the query (deployment queries do not allow use of aggregation function SUM, AVG, COUNT, MAX, MIN in the top level SELECT of a query).
- You want to use value object methods or bean methods in the query statement. For more information, see Path expressions.



- You want to interactively test an EJB query during development but do not want to repeatedly deploy your application each time you update a finder or select query.

If you have a query that has a high frequency of execution you should define it as a finder or select method and consider using SQLJ as a deployment option for best performance. The dynamic query service always uses JDBC and must parse and process the EJB query at runtime.

If you need security control over which queries a user can execute, you need to define the queries as finder or select methods and use EJB method authorization. The dynamic query service does not have fine grain security control at this time. You can control who is permitted access to the remote query bean and the local query bean, but once authorized a user can execute any valid query and return any data in the server.

The dynamic query API is a stateless session bean. Using the dynamic query API is similar to using any other J2EE EJB application bean.

The default JNDI name for the query bean is as follows, but your system administrator can change this name:

```
com/ibm/websphere/ejbquery/Query
```

The system administrator might need to install the query.ear application into the application server. The WebSphere product install does this only for the default server.

**Note:** The query.ear file is located in the <WAS\_HOME>/installableApps directory where <WAS\_HOME> is the location of the WebSphere Application Server.

The query bean has both a remote and a local interface to support both remote and local clients.

- remote interface = com.ibm.websphere.ejbquery.Query
- remote home interface = com.ibm.websphere.ejbquery.QueryHome
- local interface = com.ibm.websphere.ejbquery.QueryLocal
- local home interface = com.ibm.websphere.ejbquery.QueryLocalHome

To use the local interface or bean methods on the remote interface of the query bean, you must configure your server to use the following:

```
Application Classloader Policy = SINGLE
```

Using a value of MULTI may result in your application being unable to find the local interface for the query bean home.

The following Jar files comprise the query service:

Jar	Location	Usage
query.jar	server - AppServer/lib	query parser
qjcup.jar	server - AppServer/lib	auxillary classes for parser
querybean.jar	server - installedApps	query session bean
querymd.jar	server - AppServer/lib	auxillary classes for parser
queryws.jar	server - AppServer/lib	adapter classes for runtime
qryclient.jar	client	client stubs and classes

1. To execute a query, have your client do a JNDI lookup for the QueryHome and create an instance of the query bean. The query bean contains the executeQuery() method, which takes as parameters the query statement in the form of a string, and input parameters in the form of an array of java.lang.Object values. Remote clients also pass as arguments the size of the result set to return.

The results of the query are returned for remote clients as:

```
com.ibm.websphere.ejbquery.QueryIterator
```

or for local clients as:

```
com.ibm.websphere.ejbquery.QueryLocalIterator
```

If you want to return remote EJB references from the query, or if the query statement contains remote methods, you must use the query remote interface.

If you want to return local EJB references from the query, or if the query statement contains local methods, you must use the query local interface.

Calling the next() method on the iterator returns an instance of com.ibm.websphere.ejbquery.IQueryTuple, which contains the actual data values or object references. The iterator also contains the following methods:

```
getFieldName(int i)
getFieldsCount();
```

2. Compile and run your client program with the file **qryclient.jar** in the classpath.

For more details, see the following articles:

- Javadoc for package com.ibm.websphere.ejbquery.
- Class IQueryTuple
- Class QueryIterator

The WebSphere Sample Gallery contains a complete working example of the dynamic query service.

**Security Considerations.** WebSphere does not have security access control for CMP and CMR fields. You must therefore secure the query bean methods executeQuery(), prepareQuery(), and executePlan(), and the create() method on the query home; otherwise, any user is able to perform a dynamic query and retrieve data from your application.

## Example: Dynamic query remote client

The following is an example of the Dynamic query remote client with added explanation.

```
import com.ibm.websphere.ejbquery.QueryHome;
import com.ibm.websphere.ejbquery.Query;
import com.ibm.websphere.ejbquery.QueryIterator;
import com.ibm.websphere.ejbquery.IQueryTuple;
import com.ibm.websphere.ejbquery.QueryException;

try {
String query =
"select e.name as name , object(e) as emp from EmpBean e where e.salary < 50000";

InitialContext ic = new InitialContext();

Object obj = ic.lookup("com/ibm/websphere/ejbquery/Query");

QueryHome qh =
(QueryHome) javax.rmi.PortableRemoteObject.narrow(obj, QueryHome.class);
Query qb = qh.create();

You must specify a maximum size of the query result set. In this example 99 is used.
QueryIterator it = qb.executeQuery(query, null, null ,0, 99);
```

The iterator contains a collection of IQueryTuple objects. Each tuple contains one value of name and one value of object(e) from the query. To display the contents of the query result, use the following code:

```

while (it.hasNext()) {
IQueryTuple tuple = (IQueryTuple) it.next();
System.out.print(it.getFieldName(1));
String s = (String) tuple.getObject(1);
System.out.println(s);

System.out.println(it.getFieldName(2));
Emp e = (Emp) javax.rmi.PortableRemoteObject.narrow(tuple.getObject(2), Emp.class);
System.out.println(e.getPrimaryKey().toString());
}

```

The output from the program might look something like the following:

```

name Bob
emp 1001
name Dave
emp 298003
...

```

Finally, catch and process any exceptions. An exception might occur because of a syntax error in the query statement or a run-time processing error. The following is an example that catches and processes these exceptions:

```

} catch (QueryException qe) {
 System.out.println("Query Exception "+ qe.getMessage());
}

```

## Example: Dynamic query from local client

Establish a transaction context prior to calling `executeQuery` so that a call to `begin()` and `commit()` is needed. The `QueryLocalIterator` becomes invalid at end of a transaction. Therefore, you must use the iterator in the same transaction scope as the `executeQuery` call.

```

import com.ibm.websphere.ejbquery.QueryLocalHome;
import com.ibm.websphere.ejbquery.QueryLocal;
import com.ibm.websphere.ejbquery.QueryLocalIterator;
import com.ibm.websphere.ejbquery.IQueryTuple;
import com.ibm.websphere.ejbquery.QueryException;

String query =
"select e.name, object(e) from EmpBean e where e.salary < 50000 ";

InitialContext ic = new InitialContext();
QueryLocalHome qh = (LocalQueryHome) ic.lookup("java:comp/env/ejb/query");
QueryLocal qb = qh.create();
userTransaction.begin();
QueryLocalIterator it = qb.executeQuery(query, null, null);
while (it.hasNext()) {
IQueryTuple tuple = (IQueryTuple) it.next();
System.out.print(it.getFieldName(1));
String s = (String) tuple.getObject(1);
System.out.println(s);

System.out.println(it.getFieldName(2));
EmpLocal e = (EmpLocal) tuple.getObject(2);
System.out.println(e.getPrimaryKey().toString());
}
userTransaction.commit();

```

## Dynamic query service performance considerations

You can use the dynamic query service to build and execute queries against entity beans constructed dynamically at runtime, rather than defining them at deployment time. By using dynamic query you gain

the flexibility of queries defined at runtime and utilize the power of EJB-Query Language (QL). Apart from supporting all of the capabilities of an EJB-QL query, dynamic query adds additional functionality not available to standard static query. Two examples are the ability to select multiple data fields directly from the bean itself (static queries currently only allow one) and executing business methods directly in the query.

When used as a direct replacement for an equivalent static query, dynamic query is slower than the static variation. This slowdown is due to the need for parsing and building a plan for the query, in addition to executing it. In the static variation, these costs are paid at deploy time. In WebSphere Application Server Enterprise 5.0, expect to see 60% degradation in performance between executing a predefined static query and the equivalent dynamic query. Through optimizations implemented in WebSphere Application Server Enterprise 5.01, a cache for recently used query plans and other performance improvements, this degradation is reduced to about 25%.

You can effectively create more efficient and less resource intensive applications with dynamic query. For example, two data fields are required from the results of a query. Because a standard EJB-QL query can only select one data field, it is necessary to select the entire EJB object and extract the needed data from the returned results through data access methods, possibly traversing Container Managed Relationships (CMR) boundaries in the process. However, when using dynamic query, you can get both pieces of data directly from the query without additional CMR traversal or accessor methods. This principle is the key to evaluating whether or not dynamic query can be used for performance gain. You should review the amount of data that must be retrieved, in addition to the amount of business logic needed to retrieve it, for example, CMR traversal or accessor methods.

Using parameters in the query rather than literal values is another performance consideration. Under most circumstances, it is better to define conditional values as parameters in the query and then pass those parameters through the appropriate mechanisms. By using this method, you have a greater chance of matching a cached query plan and you eliminate the need to parse and build the plan from scratch. For example, "SELECT Object(o) FROM schemaname AS o WHERE o.fieldname LIKE foo", is more appropriately expressed as "SELECT Object(o) FROM schemaname AS o WHERE o.fieldname LIKE ?1" with the value *foo* passed as a parameter to the executeQuery method. The result is that any subsequent execution of a dynamic query structure that is the same, except for different string literal conditions, are registered as a plan cache hit affecting observed performance.

The dynamic query service is inherently handicapped from a performance viewpoint because of the necessary work of parsing the supplied queries dynamically and building the subsequent query plans. Despite this, the added functionality gained through the use of dynamic query, specifically the ability to select multiple data fields in a single query even across CMRs, creates opportunities to utilize dynamic query for the sake of performance improvement. The additional benefits of dynamic query in WebSphere Application Server Enterprise 5.01, like the use of parameters within the query, further improves the performance through caching and reusing a more general set of plans.

---

## Chapter 15. Using the internationalization service

The internationalization service adds APIs and tooling that enable J2EE applications to manage the distribution of internationalization information, or *internationalization context*, necessary to perform localizations within server-side application components. This topic summarizes the steps involved in using the internationalization service.

1. If you have an application that uses the WebSphere Version 4.0 internationalization service, migrate your application as needed.
2. Use the internationalization context API within application components to obtain or manage internationalization context. Servlet and enterprise bean business methods can use internationalization context to perform locale- and time zone-sensitive localizations. EJB client applications, and server components configured to manage internationalization context must use the internationalization context API to set the context elements scoped to their invocations.

You use the internationalization context API within Web service-enabled J2EE client programs and stateless session beans in the same manner as you would in conventional J2EE components, with one exception: Internationalization context propagated over Web services requests contains a time zone ID, whereas conventional RMI/IIOP requests propagate complete time zone information, including the raw offset, DST information, and so on.

3. Assemble internationalized applications.

Internationalization type specifies the internationalization policy applicable to a servlet or an enterprise bean and, in particular, indicates whether the application component or its hosting J2EE container will manage internationalization context. Container internationalization attributes can be specified for container-managed servlet and enterprise bean business methods. These attributes tailor a policy by indicating which context the container will scope to an invocation. Configuring internationalization policies declaratively prescribes, by means of the application's deployment descriptor, the distribution and management of context throughout an application.

**5.1+** While you edit the deployment descriptor for assembly, you can also set the internationalization type and configure any container internationalization attributes for the servlets and enterprise beans in your application.

You configure internationalization type and container internationalization attributes for Web service-enabled stateless session beans in the same manner as you do for conventional beans.

4. Manage the internationalization service. Use the administrative console to enable the service on all application servers.

By default, the service is enabled within J2EE client environments but is disabled on application servers. You must enable the service on all application servers hosting your application's servlets and enterprise beans in order to use internationalization context.

This also applies to J2EE Web service client environments and Web service-enabled enterprise beans.

5. Troubleshoot the internationalization service as needed. Use the administrative console to enable the trace service to log internationalization service messages when debugging your applications.

---

### Internationalization

An application that can present information to users according to regional cultural conventions is said to be *internationalized*. The application can be configured to interact with users from different localities in culturally appropriate ways. In an internationalized application, a user in one region sees error messages, output, and interface elements in the requested language. Date and time formats, as well as currencies, are presented appropriately for users in the specified region. A user in another region sees output in the conventional language or format for that region.

Historically, the creation of internationalized applications has been restricted to large corporations writing complex systems. Internationalization techniques have traditionally been expensive and difficult to implement, so they have been applied only to major development efforts. However, given the rise in

distributed computing and in the use of the World Wide Web, application developers have been pressured to internationalize a much wider variety of applications. This requires making internationalization techniques much more accessible to application developers.

In an application that is not internationalized, the interface that the user sees is unalterably written into the application code. On the other hand, localizing the displayed strings adds a layer of abstraction into the design of the application. Instead of simply printing an error message, an internationalized application represents the error message with some language-neutral information; in the simplest case, each error condition corresponds to a key. To print a usable error message, then, the application looks up the key in the configured message catalog. Each message catalog is a list of keys with associated strings. Different message catalogs provide strings for the different languages supported. The application looks up the key in the appropriate catalog, retrieves the corresponding error message in the requested language, and prints this string for the user.

Localization of text can be used for far more than translating error messages. For example, by using keys to represent each element in a graphical user interface (GUI) and by providing the appropriate message catalogs, the GUI itself (buttons, menus, and so on) can support multiple languages. Extending support to additional languages requires that you provide message catalogs for those languages; in many cases, the application itself needs no further modification.

Internationalization of an application is driven by two variables, the time zone and the locale. The time zone indicates how to compute the local time as an offset from a standard time like Greenwich Mean Time. The locale is a collection of information about language, currency, and the conventions for presenting information like dates. In a localized application, the locale also indicates the message catalog from which an application is to retrieve message strings. A time zone can cover many locales, and a single locale can span time zones. With both time zone and locale, the date, time, currency, and language for users in a specific region can be determined.

**5.1 +** WebSphere Business Integration Server Foundation provides an Internationalization Service, which manages the distribution of locale and time-zone information, or *internationalization context*, within Java applications.

---

## Internationalization service: Overview

In a distributed client-server environment, application processes can run on different machines, configured to different locales, corresponding to different cultural conventions; they can also be located across geographical boundaries. For an understanding of how these differences impact application development, review the topic Internationalization challenges in distributed applications.

The J2EE platform provides support for application components executing on computers with differing endian architecture and code sets, but does not provide dedicated support for application components that run on computers having different locales or time zones.

The conventional method for solving locale and time zone mismatch across remote application components is to pass one or more extra parameters on all business methods needed to convey the client-side locale or time zone to the server. Although simple, this technique has the following limitations when used in EJB applications:

- It is intrusive because it requires that one or more parameters be added to all bean methods in the call chain to locale-sensitive or time zone-sensitive methods.
- It is inherently error-prone.
- It is impracticable within applications that do not afford modification, such as legacy applications.



## The internationalization service solution

The WebSphere internationalization service addresses the challenges posed by locale and time zone mismatch without incurring the limitations of conventional techniques. It does this by systematically managing the distribution of internationalization contexts across the various components of EJB applications, including client applications, enterprise beans, and servlets.

The service works by associating an internationalization context with every thread of execution within an application. When a client-side component invokes a business method, the internationalization service interposes by obtaining the internationalization context associated with the current thread of the client-side process and attaching that context to the outgoing request. On the server-side, the internationalization service again interposes by detaching the context from the incoming request and associating it with the thread of the server-side process on which the business method will execute, effectively scoping the context to the business method. The service propagates internationalization context on subsequent business method invocations in the same manner and thus distributes the context of the originating request over the entire chain of business method invocations.

This basic operation of scoping and propagation is defined precisely by internationalization context management policies. Internationalization policies specify whether an application component or its hosting J2EE container are to manage internationalization context. For container-managed components, the policy indicates which internationalization context the container will scope to invocations on that component. Server components configured to manage internationalization context, as well as EJB clients, must use the internationalization context API to manage the internationalization context elements scoped to their invocations.

Every application component has a default policy, which can be overridden and tailored for servlets and enterprise beans at assembly time.

At execution time, application components can use the internationalization context API to get any element of the internationalization contexts scoped to an invocation. To programmatically access context elements, application components first resolve an internationalization context API reference, then invoke the appropriate API method to access the various context elements, such as the caller locale or the invocation time zone. These elements can be used in calls to Java 2 SDK internationalization API methods; for example, to perform localizations such as formatting messages, configuring dates, or comparing strings.

## Internationalization challenges in distributed applications

With the advent of Internet-based business computational models, such as eCommerce, applications increasingly comprise clients and servers that operate in different locales and geographical regions. These differences introduce challenges to the task of designing a sound client-server infrastructure. Specifically, clients and servers can:

- Execute on computers having different endian architectures or code sets.
- Be located in different locales.
- Be located in different time zones.

The following sections describe these three challenges in more detail.

### Computers with differing endian architectures or code sets

Clients and servers can reside in computers having different endian architectures: a client could reside in a little-endian CPU, while the server code runs in a big-endian one. As a more complex instance, a client might want to invoke a business method on a server running in a code set different from that of the client.

A client-server infrastructure must define precise endian and code set tracking and conversion rules. Both CORBA and J2EE have addressed the problems of endian and code set mismatches. The language-neutral CORBA formalism uses byte order indicator in all marshalled data streams to indicate the



byte order of the originating machine; in case of an endian mismatch, the receiving side can perform byte swapping for endian correction. The code set mismatch is addressed by CORBA using a comprehensive framework for code set conversion.

J2EE has nearly eliminated these problems in a unique way by relying on its Java Virtual Machine (JVM), which encodes all string data in UCS-2 format and externalizes everything in big-endian format. The JVM employs a set of platform-specific programs for interfacing with the native platform. These programs perform any necessary code set conversions between UCS-2 and the native code set of a platform.

### **Computers located in different locales**

Client and server processes can execute in geographical locations having different locales. For example, a Spanish client might invoke a business method upon an object residing on an American server. Some business methods can be locale-sensitive in nature; for example, given a business method that returns a sorted list of strings, the Spanish client expects that list to be sorted according to the Spanish collating sequence, not in the server's English collating sequence. Since data retrieval and sorting procedures run on the server, the locale of the client must be available in order to perform a legitimate sort.

A similar consideration applies in instances where the server has to return strings containing date, time, currency, exception messages, and so on, formatted according to the client's cultural expectations. Neither the CORBA nor the J2EE specifications have architecturally addressed the locale mismatch problem and other options involving extra parameters are not practical or have limitations. For example, requiring an extra parameter could require interface changes, which is a serious concern for deployed applications.

### **Computers located in different time zones**

Client and server processes can execute in geographical locations having different time zones. To date, all internationalization literature and resources have concentrated mainly on code set and locale-related issues. They have generally ignored the time zone issue, even though business methods can be sensitive to time zone as well as to locale.

For example, suppose that a vendor makes the claim that "orders received before 2:00 PM will be processed by 5:00 PM the same day". The times given, of course, are in the time zone of the server that is processing the order. It is important to know the client's time zone in order to give customers in other time zones the correct times for same-day processing.

Other time zone-sensitive operations include time stamping messages logged to a server, and accessing file or database resources. The concept of Daylight Savings Time (DST) further complicates the time zone issue. Neither the CORBA nor the J2EE specifications address time zone issues adequately, and conventional methods of solving this problem are limited.

---

## **Migrating internationalized applications**

Applications that used the internationalization service in WebSphere Application Server Version 4 can use the service in Version 5 with no modification. The packaging and structure of the internationalization context API remain identical across releases. Most importantly, the semantics of the API remain as well.

In Version 4, the internationalization service did not provide internationalization deployment descriptor policy information to direct how the service manages internationalization context across the various application components. Rather, the service employed the implicit client-side internationalization (CSI) and server-side internationalization (SSI) policies, which dictated how the service managed context according to the type of J2EE container hosting a component. For details, refer to the combined information center for WebSphere Application Server Version 4. Briefly, all server components in Version 4 were SSI, and all EJB client applications were CSI.

In Version 5, the internationalization type setting of all server components is configured to Container by default. The internationalization service assigns the default container internationalization attribute, RunAsCaller, to any container-managed (CMI) servlet or enterprise bean invocation lacking a container internationalization attribute. Hence, invocations of server components lacking internationalization policy information in the deployment descriptor run under the policy, [CMI, RunAsCaller], which is semantically equivalent to the SSI internationalization policy of Version 4; EJB client applications run under the logical policy [AMI, RunAsServer], which is equivalent to the CSI policy of Version 4.

When migrating a Version 4 application to Version 5, it is unnecessary to configure the internationalization deployment descriptor information during application assembly because all component invocations execute under semantically equivalent internationalization context management policies.

---

## Assembling internationalized applications

**5.1+** Use the Assembly Toolkit to configure internationalization in the deployment descriptors for servlets and enterprise beans.

### 1. Set the **internationalization type**.

All servlets and enterprise beans have an internationalization type setting that specifies whether internationalization context is managed by the application component or by its hosting J2EE container during invocations of their respective life cycle and business methods. Internationalization type can be configured for all server application components except entity beans, which are container-managed only.

By default, all server components employ container-managed internationalization (CMI). The default setting should suffice in most cases; when it does not, modify the internationalization type setting by completing the steps described in one of the following topics:

- “Setting the internationalization type for servlets”
- “Setting internationalization type for enterprise beans” on page 961

### 2. Set the **container internationalization attribute**.

You can associate CMI servlets, and business methods of CMI enterprise beans, with a container internationalization attribute that specifies which of three internationalization contexts - **Caller**, **Server**, or **Specified** - the container is to scope to an invocation. When running as specified, the container internationalization attribute also specifies the custom internationalization context elements.

Named container internationalization attributes can be associated with sets of servlets or with sets of EJB business methods. Initially, CMI servlets and business methods implicitly run as caller and do not associate with a container internationalization attribute. When the implicit behavior or an associated attribute setting is unsuitable, configure an attribute by completing the steps described in one of the following topics:

- “Configuring container internationalization for servlets” on page 960
- “Configuring container internationalization for enterprise beans” on page 961

## Setting the internationalization type for servlets

This task sets the internationalization type for a servlet within a Web module.

### **5.1+**

Make sure that you have enabled server target support for J2EE projects and that you set the target server as **Integration Server v5.1**.

1. **5.1+** Start the Assembly Toolkit.
2. **5.1+** If necessary, import your application.
  - a. From the menu bar, click **File > Import**.

- b. Select **EAR file** from the list.
  - c. In the Enterprise Application Import panel, enter the path for your EAR file.
  - d. Click **Finish**.
3. **5.1+** In the J2EE perspective, select the Web module for which you want to set the internationalization type.
4. **5.1+** In the Deployment Descriptor Editor, select the Extended Services tab.
5. **5.1+** From the **Servlets and JSPs** list of the Extended Services panel, select the servlet for which you want to set the internationalization type.
6. **5.1+** Under **Internationalization**, select a value from the **Internationalization type** list. Valid values are **Application managed** or **Container managed**. If you select **Container managed**, you can then set container-managed internationalization attributes for methods within the servlet. For more information, see "Configuring container internationalization for servlets."
7. **5.1+** From the menu bar, click **File > Save**.

The internationalization type setting is assigned to the servlet.

## Configuring container internationalization for servlets

This task configures container internationalization for a servlet within a Web module.

### 5.1+

Make sure that you have enabled server target support for J2EE projects and that you set the target server as **Integration Server v5.1**.

This procedure relates one or more servlets to a container internationalization attribute.

1. **5.1+** Start the Assembly Toolkit.
2. **5.1+** If necessary, import your application.
  - a. From the menu bar, click **File > Import**.
  - b. Select **EAR file** from the list.
  - c. In the Enterprise Application Import panel, enter the path for your EAR file.
  - d. Click **Finish**.
3. **5.1+** In the J2EE perspective, select the Web module for which you want to set the internationalization type.
4. **5.1+** In the Deployment Descriptor Editor, select the Extended Services tab.
5. **5.1+** From the **Servlets and JSPs** list of the Extended Services panel, select the servlet for which you want to configure container-managed internationalization (CMI) attributes. As a group, CMI attributes form an *internationalization policy*.
6. **5.1+** Under **Internationalization**, select **Container managed** from the **Internationalization type** list.
7. **5.1+** Set the **Run As** field by selecting **Caller**, **Server**, or **Specified**.
8. **5.1+** If the servlet is to be run as **Specified**, select an internationalization policy from the **Specified** list or define a new policy.
  - a. To define an internationalization policy, click **New**. The New Specified Initialization wizard opens.
  - b. In the **Description** field, give the policy a name.

- c. Set a time zone ID and add a time zone description as needed. If you do not find the appropriate time zone in the ID list, click **Customize** to define one relative to Greenwich Mean Time (GMT).
- d. If needed, create a locale for the policy. To create a locale, click **Add**; select a language and (optional) geographic region; specify a variant as needed. Add a locale description and click **OK** to finish. The new locale is added to the **Locales** list.
- e. Select a locale from the **Locales** list and click **Finish**.

Selected servlets are now configured to run under the associated internationalization settings.

## Setting internationalization type for enterprise beans

This task sets the internationalization type for an enterprise bean within an EJB module.

### 5.1+

Make sure that you have enabled server target support for J2EE projects and that you set the target server as **Integration Server v5.1**.

Container-managed internationalization (CMI) is the default type; entity beans cannot be set to application-managed internationalization (AMI).

Use this also for a stateless session bean enabled for Web services.

1. **5.1+** Start the Assembly Toolkit.
2. **5.1+** If necessary, import your application.
  - a. From the menu bar, click **File > Import**.
  - b. Select **EAR file** from the list.
  - c. In the Enterprise Application Import panel, enter the path for your EAR file.
  - d. Click **Finish**.
3. **5.1+** In the J2EE perspective, select the EJB module for which you want to set the internationalization type.
4. **5.1+** Edit the deployment descriptor for the module. With the EJB module selected, right click and then select **Open with > Deployment Descriptor Editor**. The Deployment Descriptor Editor opens.
5. **5.1+** In the Deployment Descriptor Editor, select the Internationalization tab. Any enterprise beans already configured for AMI appear in the **Internationalization type** list.
6. **5.1+** To set the internationalization type to AMI, click **Add** below the **Internationalization type** list. The Internationalization Type wizard opens. Only message-driven or session beans can be selected.
7. **5.1+** Select the beans you want to set and click **Finish** to exit the wizard.
8. **5.1+** From the menu bar, click **File > Save**.

The internationalization type is assigned to the bean.

## Configuring container internationalization for enterprise beans

This task configures container internationalization for enterprise bean business methods.

### 5.1+

Make sure that you have enabled server target support for J2EE projects and that you set the target server as **Integration Server v5.1**.

This procedure relates one or more business methods to one or more container-managed internationalization (CMI) attributes.

Use this also for stateless session beans that are enabled for Web services.

1. **5.1 +** Start the Assembly Toolkit.
2. **5.1 +** If necessary, import your application.
  - a. From the menu bar, click **File > Import**.
  - b. Select **EAR file** from the list.
  - c. In the Enterprise Application Import panel, enter the path for your EAR file.
  - d. Click **Finish**.
3. **5.1 +** In the J2EE perspective, select the EJB module for which you want to set the internationalization type.
4. **5.1 +** Edit the deployment descriptor for the module. With the EJB module selected, right click and then select **Open with > Deployment Descriptor Editor**. The Deployment Descriptor Editor opens.
5. **5.1 +** In the Deployment Descriptor Editor, select the Internationalization tab. Any methods already configured appear in the **Internationalization attributes** list.
6. **5.1 +** To configure a CMI business method, click **Add** below the **Internationalization attributes** list. The Internationalization Attributes wizard opens.
7. **5.1 +** Set the **Run As** field by selecting **Caller**, **Server**, or **Specified**. Add a meaningful description. As a group, the CMI attribute settings comprise an internationalization policy. The description appears as the policy name shown in the **Internationalization attributes** list when you are finished. If you do not provide a description, the policy name appears as `null`.

If the bean is to be run as **Specified**, complete the following steps to specify the context elements that the container will scope to bean method invocations.

  - a. Set a time zone ID and add a time zone description as needed. If you do not find the appropriate time zone in the ID list, click **Custom** to define one relative to Greenwich Mean Time (GMT).
  - b. Set a locale. Select a language and (optional) geographic region; specify a variant as needed. Add a locale description as needed and click **OK** to finish.
8. **5.1 +** Click **Next**.
9. **5.1 +** Select the beans for which you want to configure method-level internationalization attributes and click **Next**.
10. **5.1 +** Select the methods that you want to configure and click **Next**. A check box appears next to each method name that you select.
  - Clicking **Apply to All** places a check box next to the displayed bean name.
  - Clicking **Select Beans** enables you to select more beans with CMI.
11. **5.1 +** Click **Finish** to exit the wizard.
12. **5.1 +** From the menu bar, click **File > Save**.

The bean methods are now configured to run under the associated internationalization settings.

---

## Using the internationalization context API

Enterprise JavaBeans (EJB) client applications, servlets, and enterprise beans can programmatically obtain and manage internationalization context using the internationalization context API.

The `java.util` and `com.ibm.websphere.i18n.context` packages contain all classes necessary to use the internationalization service within an EJB application. Classes specific to the internationalization service reside in the file `WAS_HOME/lib/i18nctx.jar`. Before compiling application components that import internationalization service classes, add the `i18nctx.jar` file to your `CLASSPATH`.

1. Gain access to the internationalization context API.
2. Access caller locales and time zone.
3. Access invocation locales and time zone.

Each EJB application component uses the internationalization context API differently. Three code examples are provided that illustrate how to use the API within each application type. Differences in API usage, as well as other coding tips, are noted in comments that precede the relevant statement blocks.

For Web services client applications and stateless session beans, you use the API to obtain and manage internationalization context in the same manner as for EJB clients and enterprise beans, respectively.

## Gaining access to the internationalization context API

This topic describes how to access the internationalization service by resolving a reference to the internationalization context API.

### Tip:

Resolve internationalization context API references once over the lifecycle of an application component, within the initialization method of that component (for example, within the `init()` method of servlets, or within the `SetXxxContext()` method of enterprise beans).

For Web service client programs, perform the following task to resolve a reference to the internationalization context API during initialization. For stateless session beans enabled for Web services, resolve the reference in the `setSessionContext()` method.

1. Resolve a reference to the `UserInternationalization` interface by performing a lookup on the JNDI name `java:comp/websphere/UserInternationalization`. For example:

```
//-----
// Internationalization context imports.
//-----
import com.ibm.websphere.i18n.context.*;
import javax.naming.*;
...

public class MyApplication {
 ...

 //-----
 // Resolve a reference to the UserInternationalization interface.
 //-----
 InitialContext initCtx = null;
 UserInternationalization userI18n = null;
 final String UserI18nUrl = "java:comp/websphere/UserInternationalization";
 try {
 initCtx = new InitialContext();
 userI18n = (UserInternationalization)initCtx.lookup(UserI18nUrl);
 }
```

```

 }
 catch (NamingException ne) {
 // UserInternationalization URL is unavailable.
 }
}

```

If the `UserInternationalization` object is unavailable due to an anomaly or a restriction, the JNDI lookup invocation throws a `javax.naming.NameNotFoundException` that contains the `java.lang.IllegalStateException`.

2. Use the `UserInternationalization` reference to create references to the `CallerInternationalization` or `InvocationInternationalization` objects, which provide access to elements of the Caller or Invocation internationalization contexts, respectively. The `CallerInternationalization` reference can be bound to the `Internationalization` interface, only; the `InvocationInternationalization` reference can be bound to either the `Internationalization` or the `InvocationInternationalization` interfaces, depending on whether the application requires read-only or read-write access to invocation context. For example:

```

...
//-----
// Resolve references to the Internationalization and
// InvocationInternationalization interfaces.
//-----
Internationalization callerI18n = null;
InvocationInternationalization invocationI18n = null;
try {
 callerI18n = userI18n.getCallerInternationalization();
 invocationI18n = userI18n.getInvocationInternationalization();
}
catch (IllegalStateException ise) {
 // An Internationalization interface(s) is unavailable.
}

```

## Accessing caller locales and time zone

An application component must first resolve a reference to the `CallerInternationalization` object and then bind it to the `Internationalization` interface.

Every remote invocation of an application component has an associated caller internationalization context associated with the thread running that invocation. Caller context is propagated by the internationalization service and middleware to the target of a request, such as an EJB business method or servlet service method. Perform the following task to access elements of the Caller internationalization context.

This task also applies to Web service client programs and stateless session beans enabled for Web services.

1. Obtain the desired caller context elements.

```

java.util.Locale [] myLocales = null;
try {
 myLocales = callerI18n.getLocales();
}
catch (IllegalStateException ise) {
 // The Caller context is unavailable;
 // is the service started and enabled?
}
...

```

The `Internationalization` interface contains the following methods to get caller internationalization context elements:

- **Locale [] getLocales()** Returns the list of caller locales associated with the current thread.
- **Locale getLocale()** Returns the first in the list of caller locales associated with the current thread.
- **TimeZone getTimeZone()** Returns the caller `SimpleTimeZone` associated with the current thread.

The `Internationalization` interface allows read-only access to internationalization context within application components. Methods of the `Internationalization` interface are available to all EJB application components and are used in the same manner for each, but the method semantics vary



according to the component's type. For instance, when obtaining caller locale within an EJB client application, the interface returns the default locale of the host JVM; in contrast, when obtaining caller context within a servlet service method (for example, `doPost()` or `doGet()` methods), the interface returns the first locale (accept-language) propagated within the corresponding HTML request. See Internationalization context for a discussion of how the service propagates internationalization context throughout an application.

2. Use the caller context elements to localize computations under a locale or time zone of the calling process.

```
DateFormat df = DateFormat.getDateInstance(myLocale);
String localizedDate = df.getDateInstance().format(aDateInstance);
...
```

## Accessing invocation locales and time zone

An application component must first resolve a reference to the `InvocationInternationalization` object and then bind it to the `InvocationInternationalization` interface of the internationalization context API.

Every remote invocation of a servlet service or EJB business method has an invocation internationalization context associated with the thread running that invocation. Invocation context is the internationalization context under which servlet and business method implementations execute; it is propagated on subsequent invocations by the internationalization service and middleware. Perform the following task to access elements of the invocation internationalization context.

This task also applies to Web services clients and stateless session beans enabled for Web services.

1. Obtain the desired invocation context elements.

```
java.util.Locale myLocale;
try {
 myLocale = invocationI18n.getLocale();
}
catch (IllegalStateException ise) {
 // The invocation context is unavailable;
 // is the service started and enabled?
}
...
```

The `InvocationInternationalization` interface contains the following methods to both get and set invocation internationalization context elements:

- **Locale [] getLocales()**. Returns the list of invocation locales associated with the current thread.
- **Locale getLocale()**. Returns the first in the list of invocation locales associated with the current thread.
- **TimeZone getTimeZone()**. Returns the invocation `SimpleTimeZone` associated with the current thread.
- **setLocales(Locale [])**. Sets the list of invocation locales associated with the current thread to the supplied list.
- **setLocale(Locale)**. Sets the list of invocation locales associated with the current thread to a list containing the supplied locale.
- **setTimeZone(TimeZone)**. Sets the invocation time zone associated with the current thread to the supplied `SimpleTimeZone`.
- **setTimeZone(String)**. Sets invocation time zone associated with the current thread to a `SimpleTimeZone` having the supplied ID.

The `InvocationInternationalization` interface allows read and write access to invocation internationalization context within application components. However, according to internationalization context management policies, only components configured to manage internationalization context (AMI components) have write access to invocation internationalization context elements. Calls to set invocation context elements within CMI application components result in a `java.lang.IllegalStateException`. Any differences in how application components can use `InvocationInternationalization` methods are explained in Internationalization context.

2. Use the invocation context elements to localize a computation under a locale or time zone of the invoking process.

```
DateFormat df = DateFormat.getDateInstance(myLocale);
String localizedDate = df.getDateInstance().format(aDateInstance);
...
```

In the following code example, locale (en,GB) and simple time zone (GMT) transparently propagate on the call to the myBusinessMethod() method. Server-side application components, such as myEjb, can use the InvocationInternationalization interface to obtain these context elements.

```
...
//-----
// Set the invocation context under which the business method or
// servlet will execute and propagate on subsequent remote business
// method invocations.
//-----
try {
 invocationI18n.setLocale(new Locale("en", "GB"));
 invocationI18n.setTimeZone(SimpleTimeZone.getTimeZone("GMT"));
}
catch (IllegalStateException ise) {
 // Is the component CMI; is the service started and enabled?
}
myEjb.myBusinessMethod();
```

Within CMI application components, the Internationalization and InvocationInternationalization interfaces are semantically equivalent, and either of these interfaces can be used to obtain the context associated with the thread on which that component is running. For instance, both interfaces can be used to obtain the list of locales propagated to the servlet doPost() service method.

## Example: Internationalization context in an EJB client program

The following code example illustrates how to use the internationalization context API within a contained EJB client program.

This example also applies to Web service client programs.

```
//-----
// Basic Example: J2EE EJB client.
//-----
package examples.basic;

//-----
// INTERNATIONALIZATION SERVICE: Imports.
//-----
import com.ibm.websphere.il8n.context.UserInternationalization;
import com.ibm.websphere.il8n.context.Internationalization;
import com.ibm.websphere.il8n.context.InvocationInternationalization;

import javax.naming.InitialContext;
import javax.naming.Context;
import javax.naming.NamingException;
import java.util.Locale;
import java.util.SimpleTimeZone;

public class EjbClient {

 public static void main(String args[]) {

 //-----
 // INTERNATIONALIZATION SERVICE: API references.
 //-----
 UserInternationalization userI18n = null;
 Internationalization callerI18n = null;
 InvocationInternationalization invocationI18n = null;
```

```

//-----
// INTERNATIONALIZATION SERVICE: JNDI name.
//-----
final String UserI18NUrl =
 "java:comp/websphere/UserInternationalization";

//-----
// INTERNATIONALIZATION SERVICE: Resolve the API.
//-----
try {
 Context initialContext = new InitialContext();
 userI18n = (UserInternationalization)initialContext.lookup(
 UserI18NUrl);
 callerI18n = userI18n.getCallerInternationalization();
 invI18n = userI18n.getInvocationInternationalization ();
} catch (NamingException ne) {
 log("Error: Cannot resolve UserInternationalization: Exception: " + ne);
} catch (IllegalStateException ise) {
 log("Error: UserInternationalization is not available: " + ise);
}
...

//-----
// INTERNATIONALIZATION SERVICE: Set invocation context.
//
// Under Application-managed Internationalization (AMI), contained EJB
// client programs may set invocation context elements. The following
// statements associate the supplied invocation locale and time zone
// with the current thread. Subsequent remote bean method calls will
// propagate these context elements.
//-----
try {
 invocationI18n.setLocale(new Locale("fr", "FR", ""));
 invocationI18n.setTimeZone("ECT");
} catch (IllegalStateException ise) {
 log("An anomaly occurred accessing Invocation context: " + ise);
}
...

//-----
// INTERNATIONALIZATION SERVICE: Get locale and time zone.
//
// Under AMI, contained EJB client programs can get caller and
// invocation context elements associated with the current thread.
// The next four statements return the invocation locale and time zone
// associated above, and the caller locale and time zone associated
// internally by the service. Getting a caller context element within
// a contained client results in the default element of the JVM.
//-----
Locale invocationLocale = null;
SimpleTimeZone invocationTimeZone = null;
Locale callerLocale = null;
SimpleTimeZone callerTimeZone = null;
try {
 invocationLocale = invocationI18n.getLocale();
 invocationTimeZone =
 (SimpleTimeZone)invocationI18n.getTimeZone();
 callerLocale = callerI18n.getLocale();
 callerTimeZone = (SimpleTimeZone)callerI18n.getTimeZone();
} catch (IllegalStateException ise) {
 log("An anomaly occurred accessing I18n context: " + ise);
}

...
} // main

```

```

...
void log(String s) {
 System.out.println ((s == null) ? "null" : s);
}
} // EjbClient

```

## Example: Internationalization context in an EJB servlet

The following code example illustrates how to use the internationalization context API within a servlet. Note the `init()` and `doPost()` methods.

```

...
//-----
// INTERNATIONALIZATION SERVICE: Imports.
//-----
import com.ibm.websphere.i18n.context.UserInternationalization;
import com.ibm.websphere.i18n.context.Internationalization;
import com.ibm.websphere.i18n.context.InvocationInternationalization;

import javax.naming.InitialContext;
import javax.naming.Context;
import javax.naming.NamingException;
import java.util.Locale;

public class J2eeServlet extends HttpServlet {

 ...
 //-----
 // INTERNATIONALIZATION SERVICE: API references.
 //-----
 protected UserInternationalization userI18n = null;
 protected Internationalization i18n = null;
 protected InvocationInternationalization invI18n = null;

 //-----
 // INTERNATIONALIZATION SERVICE: JNDI name.
 //-----
 public static final String UserI18NUrl =
 "java:comp/websphere/UserInternationalization";

 protected Locale callerLocale = null;
 protected Locale invocationLocale = null;

 /**
 * Initialize this servlet.
 * Resolve references to the JNDI initial context and the
 * internationalization context API.
 */
 public void init() throws ServletException {

 //-----
 // INTERNATIONALIZATION SERVICE: Resolve API.
 //
 // Under Container-managed Internationalization (CMI), servlets have
 // read-only access to invocation context elements. Attempts to set these
 // elements result in an IllegalStateException.
 //
 // Suggestion: cache all internationalization context API references
 // once, during initialization, and use them throughout the servlet
 // lifecycle.
 //-----
 try {
 Context initialContext = new InitialContext();
 userI18n = (UserInternationalization)initialContext.lookup(UserI18NUrl);
 callerI18n = userI18n.getCallerInternationalization();
 invI18n = userI18n.getInvocationInternationalization();
 } catch (NamingException ne) {

```

```

 throw new ServletException("Cannot resolve UserInternationalization" + ne);
 } catch (IllegalStateException ise) {
 throw new ServletException ("Error: UserInternationalization is not
 available: " + ise);
 }
 ...
} // init

/**
 * Process incoming HTTP GET requests.
 * @param request Object that encapsulates the request to the servlet
 * @param response Object that encapsulates the response from the
 * Servlet.
 */
public void doGet(
 HttpServletRequest request,
 HttpServletResponse response)
 throws ServletException, IOException {
 doPost(request, response);
} // doGet

/**
 * Process incoming HTTP POST requests
 * @param request Object that encapsulates the request to
 * the Servlet.
 * @param response Object that encapsulates the response from
 * the Servlet.
 */
public void doPost(
 HttpServletRequest request,
 HttpServletResponse response)
 throws ServletException, IOException {

 ...
 //-----
 // INTERNATIONALIZATION SERVICE: Get caller context.
 //
 // The Internationalization service extracts the accept-languages
 // propagated in the HTTP request and associates them with the
 // current thread as a list of locales within the caller context.
 // These locales are accessible within HTTP Servlet service methods
 // using the caller internationalization object.
 //
 // If the incoming HTTP request does not contain accept languages,
 // the service associates the server's default locale. The service
 // always associates the GMT time zone.
 //
 //-----
 try {
 callerLocale = callerI18n.getLocale(); // caller locale
 // the following code enables you to get invocation locale,
 // which depends on the Internationalization policies.
 invocationLocale = invI18n.getLocale(); // invocation locale
 } catch (IllegalStateException ise) {
 log("An anomaly occurred accessing Invocation context: " + ise);
 }
 // NOTE: Browsers may propagate accept-languages that contain a
 // language code, but lack a country code, like "fr" to indicate
 // "French as spoken in France." The following code supplies a
 // default country code in such cases.
 if (callerLocale.getCountry().equals(""))
 callerLocale = AccInfoJBean.getCompleteLocale(callerLocale);

 // Use iLocale in JDK locale-sensitive operations, etc.
 ...
} // doPost

```

```

...
void log(String s) {
 System.out.println ((s == null) ? "null" : s);
}
} // CLASS J2eeServlet

```

## Example: Internationalization context in a session bean

The following code example illustrates how to perform a localized operation using the internationalization service within a session bean.

This example also applies to Web service-enabled session beans.

```

...
//-----
// INTERNATIONALIZATION SERVICE: Imports.
//-----
import com.ibm.websphere.i18n.context.UserInternationalization;
import com.ibm.websphere.i18n.context.Internationalization;
import com.ibm.websphere.i18n.context.InvocationInternationalization;

import javax.naming.InitialContext;
import javax.naming.Context;
import javax.naming.NamingException;
import java.util.Locale;

/**
 * This is a stateless Session Bean Class
 */
public class J2EESessionBean implements SessionBean {

 //-----
 // INTERNATIONALIZATION SERVICE: API references.
 //-----
 protected UserInternationalization userI18n = null;
 protected InvocationInternationalization invI18n = null;

 //-----
 // INTERNATIONALIZATION SERVICE: JNDI name.
 //-----
 public static final String UserI18NUrl =
 "java:comp/websphere/UserInternationalization";
 ...

 /**
 * Obtain the appropriate internationalization interface
 * reference in this method.
 * @param ctx javax.ejb.SessionContext
 */
 public void setSessionContext(javax.ejb.SessionContext ctx) {

 //-----
 // INTERNATIONALIZATION SERVICE: Resolve the API.
 //-----
 try {
 Context initialContext = new InitialContext();
 userI18n = (UserInternationalization)initialContext.lookup(
 UserI18NUrl);
 invI18n = userI18n.getInvocationInternationalization();
 } catch (NamingException ne) {
 log("Error: Cannot resolve UserInternationalization: Exception: " + ne);
 } catch (IllegalStateException ise) {
 log("Error: UserInternationalization is not available: " + ise);
 }
 } // setSessionContext

```

```

/**
 * Set up resource bundle using I18n Service
 */
public void setResourceBundle()
{
 Locale invLocale = null;

 //-----
 // INTERNATIONALIZATION SERVICE: Get invocation context.
 //-----
 try {
 invLocale = invI18n.getLocale();
 } catch (IllegalStateException ise) {
 log ("An anomaly occurred while accessing Invocation context: " + ise);
 }
 try {
 Resources.setResourceBundle(invLocale);
 // Class Resources provides support for retrieving messages from
 // the resource bundle(s). See Currency Exchange sample source code.
 } catch (Exception e) {
 log("Error: Exception occurred while setting resource bundle: " + e);
 }
} // setResourceBundle

/**
 * Pass message keys to get the localized texts
 * @return java.lang.String []
 * @param key java.lang.String []
 */
public String[] getMsgs(String[] key) {
 setResourceBundle();
 return Resources.getMsgs(key);
}

...
void log(String s) {
 System.out.println(((s == null) ? ";null" : s));
}
} // CLASS J2EESessionBean

```

## Internationalization context API: Programming reference

Application components programmatically manage internationalization context through the `UserInternationalization`, `Internationalization`, and `InvocationInternationalization` interfaces in the `com.ibm.websphere.i18n.context` package. The following code example introduces the internationalization context API:

```

public interface UserInternationalization {
 public Internationalization getCallerInternationalization();
 public InvocationInternationalization
 getInvocationInternationalization();
}

public interface Internationalization {
 public java.util.Locale[] getLocales();
 public java.util.Locale getLocale();
 public java.util.TimeZone getTimeZone();
}

public interface InvocationInternationalization
 extends Internationalization {
 public void setLocales(java.util.Locale[] locales);
 public void setLocale(java.util.Locale jmLocale);
 public void setTimeZone(java.util.TimeZone timeZone);
 public void setTimeZone(String timeZoneId);
}

```



## UserInternationalization interface

The UserInternationalization interface provides factory methods for obtaining references to the CallerInternationalization and InvocationInternationalization context objects. Use these references to access elements of the caller and invocation contexts correlated to the current thread.

Methods of the UserInternationalization interface:

### Internationalization **getCallerInternationalization()**

Returns a reference implementing the Internationalization interface that allows access to elements of the caller internationalization context correlated to the current thread. If the service is disabled, this method throws an `IllegalStateException`.

### InvocationInternationalization **getInvocationInternationalization()**

Returns a reference implementing the InvocationInternationalization interface. If the service is disabled, this method throws an `IllegalStateException`.

## Internationalization interface

The Internationalization interface declares methods affording read-only access to internationalization context. Given a caller or invocation internationalization context object created with the UserInternationalization interface, bind the object to the Internationalization interface in order to get elements of that context type. Observe that caller internationalization context can be accessed only through this interface.

Methods of the Internationalization interface:

### **Locale[] getLocales()**

Returns the chain of locales within the internationalization context (object) bound to the interface, provided the chain is not null; otherwise this method returns a chain of length(1) containing the default locale of the JVM.

### **Locale getLocale()**

Returns the first in the chain of locales within the internationalization context (object) bound to the interface, provided the chain is not null; otherwise this method returns the default locale of the JVM.

### **TimeZone getTimeZone()**

Returns the caller time zone (that is, the `SimpleTimeZone`) associated with the current thread, provided the time zone is non-null; otherwise this method returns the process time zone.

## InvocationInternationalization interface

The InvocationInternationalization interface declares methods affording read and write access to InvocationInternationalization context. Given an invocation internationalization context object created with the UserInternationalization interface, bind the object to the InvocationInternationalization interface in order to get and set elements of the invocation context.

**Note:** According to the container-managed internationalization (CMI) policy, all set methods, `setXxx()`, throw an `IllegalStateException` when called within a CMI servlet or enterprise bean.

Methods of the InvocationInternationalization interface:

### **void setLocales(java.util.Locale[] locales)**

Sets the chain of locales to the supplied chain, *locales*, within the invocation internationalization context. The supplied chain can be null or have length( $\geq 0$ ). When the supplied chain is null or has length(0), the service sets the chain of invocation locales to an array of length(1) containing the default locale of the JVM. Null entries can exist within the supplied locale list, for which the service substitutes the default locale of the JVM on remote invocations.

### **void setLocale(java.util.Locale locale)**

Sets the chain of locales within the invocation internationalization context to an array of length(1)

containing the supplied locale, *locale*. The supplied locale can be null, in which case the service instead sets the chain to an array of length(1) containing the default locale of the JVM.

**void setTimeZone(java.util.TimeZone timeZone)**

Sets the time zone within the invocation internationalization context to the supplied time zone, *time zone*. If the supplied time zone is not an exact instance of `java.util.SimpleTimeZone` or is null, the service instead sets the invocation time zone to the default time zone of the JVM.

**void setTimeZone(String timeZoneId)**

Sets the time zone within the invocation internationalization context to the `java.util.SimpleTimeZone` having the supplied ID, *timeZoneId*. If the supplied time zone ID is null or invalid (that is, it does not appear in the list of IDs returned by the `java.util.TimeZone.getAvailableIds()` method) the service sets the invocation time zone to the simple time zone having an ID of GMT, an offset of 00:00, and otherwise invalid fields.

## Internationalization context

An *internationalization context* is a distributable collection of internationalization information containing an ordered list, or chain, of locales and a single time zone, where the locales and time zone are instances of Java SDK types, `java.util.Locale` and `java.util.TimeZone`. A locale chain is ordered according to the user's preference.

The internationalization service manages and makes available two varieties of internationalization context: the *caller context*, which represents the caller's localization environment, and the *invocation context*, which represents the localization environment under which a business method runs. Server application components use elements of the caller and invocation internationalization contexts to appropriately tailor locale-sensitive and time zone-sensitive computations.

**Note:** The internationalization service does not support time zone types other than Java SDK type `java.util.SimpleTimeZone`. Unsupported time zone types silently map to the default time zone of the JVM when supplied to internationalization context API methods. For a complete description of the `java.util.Locale`, `java.util.TimeZone`, and `java.util.SimpleTimeZone` types, refer the Java SDK API documentation.

### Caller context

Caller internationalization context contains the locale chain and time zone received on incoming EJB business method and servlet service method invocations; it is the internationalization context propagated from the calling process. Use caller context elements within server application components to localize computations to the calling component. Caller context is read-only and can be accessed by all application components by using the **Internationalization** interface of the internationalization context API.

Caller context is computed in the following manner: On an EJB business method or servlet service method invocation, the internationalization service extracts the internationalization context from the incoming request and scopes this context to the method as the caller context. For any missing or null context element, the service inserts the corresponding default element of the JVM (for example, `java.util.Locale.getDefault()` or `java.util.TimeZone.getDefault()`.)

The service performs a similar insertion whenever missing or null Caller context elements are encountered on invocations of stateless session beans that are enabled for Web services.

Formally, caller context is the invocation context of the calling business method or application component.

### Invocation context

Invocation internationalization context contains the locale chain and time zone under which EJB business methods and servlet service methods execute. It is managed by either the hosting container or the application component, depending on the applicable internationalization policy. On outgoing business

method requests, it is the context that propagates to the target process. Use invocation context elements to localize computations under the specified settings of the current application component.

Invocation context is computed in the following manner: On an incoming business method or servlet service method invocation, the internationalization service queries the associated context management policy. If the policy is container-managed internationalization (CMI), the container scopes the context designated by the policy to the invocation; otherwise the policy is application-managed internationalization (AMI), and the container scopes a vacuous context to the invocation that can be altered by the method implementation.

Application components can access invocation context elements through both the **Internationalization** and **InvocationInternationalization** interfaces of the internationalization context API. Invocation context elements can be set (overwritten) under the application-managed internationalization policy only.

On an outgoing business method request, the service obtains the currently scoped invocation context and attaches it to the request. This outgoing exported context becomes the caller context of the target invocation. When supplying invocation context elements, either for export on outgoing requests or through the API, the internationalization service always provides the most recent element set using the API; also, the service supplies the corresponding default element of the JVM for any null invocation context element.

Because the internationalization context that is propagated over Web services (SOAP) requests contains a time zone ID rather than the entire state of a `java.lang.SimpleTimeZone` object, time zone information might be lost when a Web service-enabled client program or session bean becomes involved in remote business computation.

### **Internationalization context: Propagation and scope**

The scope of internationalization context is implicit. Every EJB client application, servlet service method, and EJB business method invocation has two internationalization contexts under which it executes. For each application component invocation, the container enters the caller context and the invocation context, as indicated by the pertinent internationalization policy, into scope before the container delegates to the actual implementation. When the implementation returns, the service removes these contexts from scope. The internationalization service supplies no programmatic mechanism for components to explicitly manage the scope of internationalization context.

The service scopes internationalization context differently with respect to application component type:

- EJB client programs (contained)
- Servlets
- Enterprise beans
- Web service client programs (contained)
- Stateless session beans that are enabled for Web services

Internationalization context observes by-value semantics over remote method requests, meaning that changes to internationalization context elements scoped to an invocation do not affect the corresponding elements of the internationalization context scoped to the remote calling process. Also, modifications to context elements obtained using the internationalization context API do not affect the corresponding elements scoped to the invocation.

### **EJB client programs (contained)**

Before it invokes the `main()` method of a client program, the J2EE client container introduces into scope invocation and caller internationalization contexts containing null elements. These contexts remain in scope throughout the life of program. EJB client programs are the base in a chain of remote business method invocations and, technically, do not have a logical caller context. Accessing a caller context element yields the corresponding default element of the client JVM. On outgoing EJB business method requests, the internationalization service propagates the invocation context to the target process. Any

unset (null) invocation context elements are replaced with the default of the JVM when exported by the internationalization context API or by outgoing requests.

**Tip:**

To propagate values other than the JVM defaults to remote business methods, EJB client programs, as well as AMI servlets or enterprise beans, must set (override) elements of the invocation context. To learn how to set invocation context elements, see the topic *Accessing invocation locales and time zone*.

**Servlets**

On every servlet service method (`doGet()`, `doPost()`) invocation, the J2EE web container introduces caller and invocation internationalization contexts into scope before delegating to the service method implementation. The caller context contains the accept-languages propagated in the HTTP servlet request, typically from a Web browser. The invocation context contains whichever context is indicated by the container internationalization attribute of the internationalization policy associated with the servlet. Any unset (null) invocation context elements are replaced with the default of the server JVM when exported by the internationalization context API or by outgoing requests. The caller and invocation contexts remain effective until immediately after the implementation returns, at which time the container removes them from scope.

**Enterprise beans**

On every EJB business method invocation, the J2EE EJB container introduces caller and invocation internationalization contexts into scope before delegating to the business method implementation. The caller context contains the internationalization context elements imported from the incoming IIOP request; if the incoming request lacks a particular internationalization context element, the container scopes a null element. The invocation context contains whichever context is indicated by the container internationalization attribute of the internationalization policy associated with the business method. On outgoing EJB business method requests, the service propagates the invocation context to the target process. Any unset (null) invocation context elements are replaced with the default of the server JVM when exported by the internationalization context API or by outgoing requests. The caller and invocation contexts remain effective until immediately after the implementation returns, at which time the container removes them from scope.

Consider a simple WebSphere EJB application having a Java client that invokes remote bean method, `myBeanMethod()`. On the client side, the application could use the Internationalization Service API to set invocation context elements. When the client calls `myBeanMethod()`, the service exports the client's invocation context to the outgoing request. On the server side, the service detaches the imported context from the incoming request and scopes it to the method as its caller context; it also scopes the invocation context to the method as indicated by the associated internationalization context management policy. The EJB container then calls the `myBeanMethod()`, which can use the internationalization context API to access elements of either the caller or invocation contexts. When `myBeanMethod()` returns, the EJB container removes these contexts from scope.

**Web service client programs (contained)**

Before it calls the `main()` method of a Web service client program, the J2EE client container introduces into scope both invocation and caller internationalization contexts that contain null elements. These contexts remain in scope throughout the duration of program execution. Web service client programs are the base in a chain of remote business method invocations and, technically, do not have a logical caller context. Accessing a Caller context element yields the corresponding default element of the client virtual machine.

On outgoing Web service requests, the internationalization service transparently creates a Simple Object Access Protocol (SOAP) header block that contains the invocation context associated with the current thread; the SOAP representation of invocation context is propagated through the request to the target process. Any unset (that is, null) invocation context elements are replaced with the default element of the JVM when exported by the internationalization context API or by outgoing requests. Also, because the header contains only a time zone ID, the additional state of the invocation context's time zone object (`java.lang.SimpleTimeZone`) might be lost, because it does not get propagated through the request.

**Tip:**

To propagate values other than the JVM defaults to remote business methods, Web service client programs, as well as AMI servlets or enterprise beans, must set (override) elements of the invocation context. For more information, see "Accessing invocation locales and time zone."

**Stateless session beans that are enabled for Web services**

On every method invocation of a Web service-enabled bean, the EJB container introduces caller and invocation internationalization contexts into scope before delegating control to the business method implementation. The caller context contains the internationalization context elements imported from the SOAP header block of the incoming request. If the incoming request lacks a particular internationalization context element, the container introduces a null element into scope. The invocation context contains whichever context is indicated by the container internationalization attribute of the internationalization policy that is associated with the business method.

On outgoing EJB business method requests, the service propagates the invocation context to the target process. Any unset (that is, null) invocation context elements are replaced with the default element of the server JVM when exported by the internationalization context API or by outgoing requests. The caller and invocation contexts remain effective until immediately after control returns from the business method implementation, at which time the container removes them from scope.

On outgoing Web service requests, the internationalization service transparently creates a SOAP header block that contains the invocation context associated with the current thread. The SOAP representation of the invocation context is propagated through the request to the target process. Any unset (that is, null) invocation context elements are replaced with the default element of the JVM when exported by the internationalization context API or by outgoing requests.

**Thread association considerations**

The Web and EJB containers scope internationalization contexts to a method by associating it with the thread that executes the method's implementation. Similarly, methods of the internationalization context API either associate context with, or obtain context associated with, the thread on which these methods execute. In cases where new threads are spawned within an application component (for instance, a user-generated thread inside the `service()` method of a servlet, or a system-generated event handling thread in an AWT client) the internationalization contexts associated with the parent thread does not automatically transfer to the newly-spawned thread. In such instances, the service exports the default locale and time zone of the JVM on any remote business method request and on any API calls executed on the new thread. If the default context is inappropriate, the desired invocation context elements must be explicitly associated to the new thread using the `setXxx()` methods of the **InvocationInternationalization** interface. Currently, internationalization context management policies allow invocation context to be set within EJB client programs, as well as within servlets, session beans, and message-driven beans employing application-managed internationalization.

**Example: Internationalization context in a SOAP header:** The following code example illustrates how internationalization context is represented within the Simple Object Access Protocol (SOAP) header of a Web service request.

```

<InternationalizationContext>
 <Locales>
 <Locale>
 <LanguageCode>ja</LanguageCode>
 <CountryCode>JP</CountryCode>
 <VariantCode>Nihonbushi</VariantCode>
 </Locale>
 <Locale>
 <LanguageCode>fr</LanguageCode>
 <CountryCode>FR</CountryCode>
 </Locale>
 <Locale>
 <LanguageCode>en</LanguageCode>
 <CountryCode>US</CountryCode>
 </Locale>
 </Locales>
 <TimeZoneID>JST</TimeZoneID>
</InternationalizationContext>

```

## Internationalization context: Management policies

Internationalization policies declaratively prescribe how J2EE application components or their hosting containers (the service) will manage internationalization context on component invocations. There are two internationalization context management policies applicable to all component types:

- Application-managed internationalization (AMI)
- Container-managed internationalization (CMI)

These policies are represented in two parts:

- Internationalization type
- Container internationalization attribute

The service defines a default, or implicit, internationalization policy for every application component type. At development time, assemblers can override the default policy for server component types by explicitly configuring their internationalization type, and optional container internationalization attributes. Policies configured during assembly are preserved in the application's deployment descriptor.

All components have an internationalization type that indicates whether it is AMI or CMI; that is, whether a component is to deploy under the application-managed or the container-managed internationalization policy. Application assemblers can set the internationalization type for servlets, session beans, and message-driven beans. Entity beans are implicitly CMI and EJB clients are implicitly AMI; neither can be configured otherwise.

For CMI servlets and enterprise beans, optional container internationalization attributes can be specified to indicate which invocation internationalization context the container is to scope to service or business methods. A CMI service or business method invocation can run under the context of the caller's process, under the default context of the server's JVM, or under a custom context specified in the attribute. Assemblers can specify one container internationalization attribute per disjoint set of CMI servlets within a Web module, or one Attribute per disjoint set of business methods of CMI beans within an EJB module. In other words, a container internationalization attribute can be associated with more than one method, but a method cannot be associated with more than one attribute.

When a WebSphere application server launches an application, the internationalization service collects policy information from the deployment descriptor, then uses this information to construct and associate an internationalization policy to every component invocation. A policy is denoted as:

```
[<Internationalization Type>,<Container Internationalization Attribute>]
```

There are several cases where the deployment descriptor appears to lack policy information, for example: EJB client applications have no configurable internationalization policy settings; AMI components do not have container internationalization attributes; and you are not required to specify container internationalization attributes for CMI components. When the service cannot obtain the explicit



internationalization type and container attribute settings from a well-formed deployment descriptor, it implicitly inserts the appropriate setting into the policy.

The service observes the following conventions when applying policies to invocations:

- Servlets (service) and EJB business methods lacking all internationalization policy information in the deployment descriptor implicitly execute under policy [CMI,RunAsCaller].
- CMI servlets and business methods lacking a container internationalization attribute in the deployment descriptor implicitly execute under policy [CMI,RunAsCaller].
- AMI servlets and business methods always lack container internationalization attributes in the deployment descriptor, but implicitly execute under the logical policy [AMI,RunAsServer].
- EJB clients always lack internationalization policy information in the deployment descriptor. By definition, EJB clients are implicitly AMI and run under the invocation context of the JVM; they execute under the logical policy [AMI,RunAsServer].

For conditions other than these, such as a malformed deployment descriptor, refer to the topic Internationalization service errors.

Internationalization policies for EJB clients and HTTP clients cannot be configured; HTTP clients do, however, run under the language priority settings of the hosting Web browser. These settings are configurable under the options dialog of most Web browsers; refer to your Web browser's documentation for details.

### ***Internationalization type:***

Every server application component has an *internationalization type* setting that indicates whether the invocation internationalization context is to be managed by the component or by the hosting J2EE container.

Server application components can be deployed to use one of two types of internationalization context management:

- Application-managed internationalization (AMI)
- Container-managed internationalization (CMI)

A server component may be deployed as AMI or CMI, but not both; CMI is the default. The setting applies to the entire component on every invocation. Entity beans use CMI and cannot be configured otherwise. EJB client applications do not have an internationalization type setting; they implicitly use AMI.

### **Application-managed internationalization (AMI)**

Under the AMI deployment policy, component developers assume complete control over the invocation internationalization context. AMI components can use the internationalization context API to programmatically set invocation context elements.

AMI components are expected to manage invocation context. Invocations of AMI components implicitly run under the default locale and time zone of the hosting JVM. Invocation context elements not set using the API default to the corresponding elements of the JVM when accessed through the API or when exported on business methods. To export context elements other than the JVM defaults, AMI servlets, AMI enterprise beans, and EJB client applications must set (overwrite) invocation elements using the internationalization context API. Moreover, the container logically suspends caller context imported on AMI servlet lifecycle method and AMI EJB business method invocations. To continue propagating the context of the calling process, AMI servlets and enterprise beans must use the API to transfer caller context elements to the invocation context.

Specify AMI for server components that have internationalization context management requirements not supported by container-managed internationalization (CMI).



## Container-managed internationalization (CMI)

CMI is the preferred internationalization context management policy for server application components; it is also the default policy. Under CMI, the internationalization service collaborates with the Web and EJB containers to set the invocation internationalization context for servlets and enterprise beans. The service sets invocation context according to the container internationalization attribute of the policy associated with a servlet (service method) or an EJB business method.

A CMI policy contains a container internationalization attribute that indicates which internationalization context the container is to scope to an invocation. For details, see topic Container internationalization attributes. By default, invocations of CMI components run under the caller's internationalization context; or rather, they adhere to the implicit policy `[CMI,RunasCaller]` whenever the servlet or business is not associated with an attribute in the deployment descriptor. For complete details, see the topic Internationalization context: Management policies.

Methods within CMI components can obtain elements of the invocation context using the internationalization context API, but cannot set them. Any attempt to set invocation context elements within CMI components results in a `java.lang.IllegalStateException`.

Specify container-managed internationalization for server application components requiring standard internationalization context management, then specify the container internationalization attributes for CMI servlets and for business methods of CMI enterprise beans that should not run under the caller's internationalization context.

### **Container internationalization attributes:**

The internationalization policy of every CMI servlet and EJB business method has a *container internationalization attribute* that specifies which internationalization context the container is to scope to its invocation.

The container internationalization attribute has three main fields:

- Run as
- Locales
- Time zone ID

As a convenience, you can create named container internationalization attributes and associate them to the following:

- Subsets of CMI servlets within a Web module
- Subsets of business methods of CMI enterprise beans within an Enterprise JavaBeans (EJB) module
- Subsets of business methods of Web service-enabled session beans. In the following descriptions, the term *supported enterprise bean* refers to both CMI enterprise beans and Web service-enabled session beans.

### **Run-as field**

The **Run-as** field specifies one of three types of invocation context that a container can scope to a method. For servlet service and EJB business methods, the container constructs the invocation internationalization context according to the **Run as** field setting and associates this context to the current thread before delegating to the method's implementation.

By default, invocations of servlet service methods and EJB business methods implicitly run as caller (`RunAsCaller`) unless the **Run as** field of a policy's attribute specifies otherwise. EJB client applications and AMI server components always run as server (`RunAsServer`).

Invocation context types specifiable with the **Run as** field are:

**Caller** The container invokes the method under the internationalization context of the calling process. For

any missing context element, the container supplies the corresponding default context element of the JVM. Select `run as caller` when you want the invocation to execute under the invocation context of the calling process.

#### **Server**

The container invokes the method under the default locale and time zone of the JVM. Select `run as server` when you want the invocation to execute under the invocation context of the JVM.

#### **Specified**

The container invokes the method under the internationalization context specified in the attribute. Select `run as specified` when you want the invocation to execute under the custom invocation context specified in the policy, then provide the custom context elements by completing the `Locales` and `Time zone ID` fields.

**Note:** JMS messages do not contain internationalization context. Although container-managed message-driven beans can be configured to run as caller, the container associates the default elements of the server process when invoking the `onMessage()` method of any message-driven bean configured as `[CMI, RunAsCaller]`.

You can also configure the **Run as** field for Web service business methods.

#### **Locales field**

The **Locales** field specifies an ordered list of locales that the container scopes to an invocation. A locale represents a specific geographical, cultural, or political region and contains three fields:

- **Language code.** Ideally, language code is one of the lower-case, two-character codes defined by ISO-639; however, language code is not restricted to ISO codes and is not a required field. A valid locale must specify a language code if it does not specify a country code.
- **Country code.** Ideally, country code is one of the upper-case, two-character codes defined by ISO-3166; however, country code is not restricted to ISO codes and is not a required field. A valid locale must specify a country code if it does not specify a language code.
- **Variant.** Variant is a vendor-specific code. Variant is not a required field and serves only to supplement the language and country code fields according to application- or platform-specific requirements.

A valid locale must specify at least a language code or a country code; the variant is always optional. The first locale of the list is returned when accessing invocation context using the internationalization context API method `getLocale()`.

#### **Time zone ID field**

The **Time zone ID** field specifies a shorthand identifier for a time zone that the container scopes to an invocation.

You can also configure the **Time zone ID** field for Web service business methods.

A time zone represents a temporal offset and computes daylight savings information. A valid ID indicates any time zone supported by the SDK type, `java.util.TimeZone`. Specifically, a valid ID is any of the IDs appearing in the list of time zone IDs returned by method `java.util.TimeZone.getAvailableIds()`, or a custom ID having the form `GMT[+|-]hh[[:]mm]`; for example, `America/Los_Angeles`, `GMT-08:00` are valid time zone IDs.

---

## **Managing the internationalization service**

To use internationalization context in an Enterprise JavaBeans (EJB) application, the internationalization service must be enabled in the run-time environments for all server-side components (servlets and enterprise beans) as well as all client-side components (EJB client applications).

During installation of this product, if you installed the **Additional services** option but do not require the internationalization service, disable the service on all J2EE clients and application servers. Disabling the service eliminates any possible performance degradation incurred by the implicit distribution of internationalization resources.

**Note:** The internationalization service cannot be enabled for HTTP clients because support for internationalization in that case is provided by the browser, not by the application server.

1. As needed, enable or disable the internationalization service for servlets and enterprise beans. The service is disabled by default within WebSphere enterprise application servers. You enable the service by using either the administrative console or the wsadmin tool.

This step applies also to Web service-enabled session beans.

2. As needed, enable or disable the internationalization service for EJB clients within the hosting WebSphere J2EE client environments. The service is enabled by default within the WebSphere J2EE client container.

This step applies also to Web service clients.

## Enabling the internationalization service for servlets and enterprise beans

Any servlet or enterprise bean can use internationalization context if the internationalization service is enabled within the hosting WebSphere application server instance.

This task applies also to Web service-enabled enterprise beans.

1. Start the administrative console.
2. Select **Servers > Application Servers > server\_name > Internationalization Service**.
3. Enable the internationalization service.
  - a. If not already selected, select the **Startup** checkbox.
  - b. Click **OK**.

When the **Startup** setting is selected, the application server automatically initializes, starts, and enables the internationalization service whenever the server starts. If you change this setting, be sure to restart the application server in order for the new setting to take effect.

To disable the service, clear the **Startup** checkbox. In this case, the internationalization service is initialized but is neither started nor enabled when the application server starts.

### Administration through scripting

Alternatively, the internationalization service can be enabled from the command line by using the wsadmin tool. To do this, start the wsadmin tool and enter the following commands:

```
set x [$AdminConfig list I18NService]
$AdminConfig modify $x { { enable true } }
$AdminConfig save
exit
```

If you enable or disable the internationalization service, be sure to stop and then restart the application server in order for the new setting to take effect.

### Internationalization service settings

Use this page to enable or disable the internationalization service.

The internationalization service manages the implicit propagation and scoping of locale and time zone information, called *internationalization context*, within enterprise applications. When the service is enabled,

server-side application components can use the internationalization context API to programmatically manage locale and time zone information, or to use this information with the J2SE Internationalization API to perform server-side localizations.

To view this administrative console page, click **Servers > Application Servers > *server\_name* > Internationalization Service** .

### **Startup:**

Specifies whether the server will attempt to start the internationalization service.

<b>Default</b>	Selected
<b>Range</b>	Valid values are Selected or Cleared

More information about valid values follows:

#### **Selected**

When the application server starts, it attempts to start the internationalization service automatically.

#### **Cleared**

The server does not try to start the internationalization service. If internationalization is to be used in applications that run on this server, the system administrator must select this property then restart the server.

## **Enabling the internationalization service for EJB clients**

By default, the internationalization service is enabled for use within Enterprise JavaBeans (EJB) client applications whenever the **i18nctx.jar** file is in the CLASSPATH setting constructed by the launchClient tool.

The internationalization service is also enabled for Web service-enabled clients.

When invoking a Java client application, the launchClient tool sets the CLASSPATH to include the i18nctx.jar file and then activates the WebSphere J2EE client (container), which initializes, starts, and enables the service before delegating to the specified application.

To disable the service for all application server instances in your installation, remove the i18nctx.jar file from the *install\_root/lib* directory. This prevents the file from inadvertently being included in the CLASSPATH setting constructed by the launchClient tool.

To selectively disable the service, include the argument `-CCDI18NService.enable=false` or `-CCDI18NService.enable=no` when invoking the launchClient tool.

---

## **Troubleshooting the internationalization service**

To have your application server emit trace statements for the internationalization service, specify the appropriate trace string to the server's diagnostic trace service.

1. Start the administrative console.
2. Select **Servers > application servers > *server\_name* > Diagnostic Trace Service**.
3. Select the **Enable Trace** checkbox.
4. In the **Trace Specification** field, type the following as a continuous string (no spaces and no line breaks):

```
com.ibm.ws.i18n.context.*=all=enabled:
com.ibm.websphere.i18n.context.*=all=enabled
```
5. Click **OK**.

6. Click **Save** on the taskbar.
7. Click **Save** in the **Save to Master Repository** panel.

These settings enable the internationalization service trace when you start or restart the corresponding application server.

## Internationalization service errors

The following conditions can occur while your internationalized application is running. These conditions might cause the internationalization service not to start, to throw instances of `IllegalStateException`, or to exercise default behaviors:

- The service is disabled
- The service is not started
- Invalid context element
- Missing context element
- Invalid policy
- Missing policy

If you encounter unexpected or exceptional behavior, the problem is likely related to one of these conditions. You need to examine the trace log to investigate these conditions, which requires that you configure the diagnostic trace service to generate messages about internationalization service function. To do this, see the topic [Troubleshooting the internationalization service](#).

### The service is disabled

The internationalization service does not initialize and start when the service's startup setting is cleared. The service generates a message indicating whether it is enabled or disabled. Applications cannot access the internationalization API when the service is disabled. If an application attempts a JNDI lookup to obtain the `UserInternationalization` reference, the lookup fails with a `NamingException` indicating the reference could not be found. In addition, the service does not scope (propagate) internationalization context on incoming (outgoing) business method invocations.

### The service is not started

The internationalization service is operational whenever it is in the `STARTED` state. For example, if an application attempts to access internationalization context and the service is not started, the API throws an `IllegalStateException`. In addition, the service does not provide runtime support for servlets and enterprise beans.

As an application server progresses through its lifecycle, it initializes, starts, stops, and terminates (destroys) the internationalization service. If an anomaly occurs during initialization, the service does not start. Once the service has been started, its state can change to `BLOCKED` in the event that a serious error occurs. The service generates a message for every state change.

If a trace message indicates that the service is not `STARTED`, examine previous messages to determine the problem. For instance, the internationalization service does not start if the activity service is unavailable and a message is displayed to that effect during initialization of the internationalization service.

During startup, the following messages indicate potential configuration or run-time problems:

#### **No ORB support**

The service could not obtain an instance of the ORB. This is a fatal error. Examine the logs for information.

#### **No TCM support**

The service could not obtain an instance of its thread context manager. This is a fatal error. Examine the logs for information.

**No IOP (Activity service) support**

The service could not register with the Activity service. This is a fatal error. The internationalization service cannot propagate or receive context on IOP requests without Activity service support. Review the logs for error conditions related to the Activity service.

**No AsynchBeans support**

The service could not register into the AsynchBeans environment. This warning indicates that the AsynchBeans environment cannot support internationalization context. If the application server should have AsynchBeans support, verify that the `asynchbeans.jar` and `asynchbeansimpl.jar` files exist in the classpath and review the trace log for any AsynchBeans error conditions.

**No EJB container support**

The service could not register with the EJB container. This is a warning that the internationalization service cannot support enterprise beans. Without EJB container support, internationalization contexts do not scope properly to EJB business methods. Review the trace log for any EJB container-related error conditions.

**No Web container support**

The service could not register with the Web container. This is a warning that the internationalization service cannot support servlets and Java Server Pages (JSPs). Without Web container support, internationalization contexts do not scope properly to servlet service methods. Review the trace log for any Web container-related error conditions.

**No Meta-data support**

The service could not register with the meta-data service. This is a warning that the internationalization service cannot process the internationalization policies within application deployment descriptors. Without meta-data support, the service associates the default internationalization context management policy, [CMI, RunAsCaller], to every servlet lifecycle method and enterprise bean business method invocation. Review the trace log for any meta-data service-related error conditions.

**No JNDI (Name service) support**

The service could not bind the `UserInternationalization` object into the namespace. This is a fatal error. Application components are unable to access internationalization context API references, and are therefore unable to access internationalization context elements. Review the trace log for any Naming (JNDI) service-related error conditions.

**No API support**

The service could not obtain an instance of an internationalization context API object. This is a fatal error. Application components are unable to access internationalization context API references, and are therefore unable to access internationalization context elements.

**Invalid context element**

The service detected an invalid internationalization context element. For example, the internationalization service does not support `TimeZone` instances of a type other than `java.util.SimpleTimeZone`. If the service encounters an invalid element, it logs a message and substitutes the corresponding default element of the JVM.

**Missing context element**

The service detected a missing internationalization context element. Incoming requests, for example from application servers not supporting the internationalization service will lack internationalization context. When the service attempts to access a caller internationalization context element, which does not exist in this case, it logs a message and substitutes the corresponding default element of the JVM.

Whenever possible, the internationalization service should be enabled within all clients and hosting application servers comprising a WebSphere enterprise application. For more information see the topic [Managing the internationalization service](#).

## Invalid policy

The internationalization service detected a malformed internationalization policy in the application deployment descriptor. At execution, the service replaces the malformed attribute with the appropriate default. For instance, if the internationalization type for an entity bean is set to **Application** during the execution of a servlet or EJB business method invocation, the service logs the inconsistency and enforces the **Container** setting instead.

Also, AMI application components do have an implicit container internationalization attribute. By default they run as server. The service silently enforces the implicit policy, [AMI, RunAsServer], and logs messages to this effect.

Invalid container internationalization attributes are likely to occur when specifying the Locales and Time zone ID fields. When encountering invalid Locales and Time zone ID within attributes, the service replaces each with the corresponding default element of the JVM. Be sure to follow the guidelines provided in the topic Assembling internationalized applications.

## Missing policy

The service detected a missing internationalization policy. At execution, the service replaces the missing policy with the appropriate default. For instance, if the internationalization type is missing for a servlet or enterprise bean, the service sets the attribute to Container.

Container internationalization attributes are not mandatory for CMI application components. In the event that a CMI servlet or EJB business method lacks a container internationalization attribute, the service silently enforces the implicit policy [CMI, RunAsCaller].

When an application lacks internationalization policies in its deployment descriptor, or meta-data support is unavailable, the service logs a message and applies the policy [CMI, RunAsCaller] on every servlet service method and EJB business method invocation.

For more information, see the following topics:

- Assembling internationalized applications
- Container internationalization attributes
- Internationalization type
- Migrating internationalized applications

## Internationalization service exceptions

The internationalization service employs one exception: `java.lang.IllegalStateException`. This exception indicates one of the following things:

- An application component attempted an operation not supported by the service's programming model. `IllegalStateException` is thrown whenever a server application component whose internationalization type is set to container-managed Internationalization (CMI) attempts to set invocation context. This is a violation of the CMI policy, under which servlets and enterprise beans cannot modify their invocation internationalization context.
- An anomaly occurred that disabled the service.

For instance, if the internationalization service does not properly initialize, the JNDI lookup on the `UserInternationalization` URL throws a `javax.naming.NameNotFoundException` containing an instance of `IllegalStateException`. Refer to the trace log to determine the reason for failure and, if necessary, contact your IBM support representative.



---

## Internationalization: Resources for learning

Use the following links to find relevant supplemental information about internationalization. The information resides on IBM and non-IBM Internet sites, whose sponsors control the technical accuracy of the information.

These links are provided for convenience. Often, the information is not specific to this product but is useful all or in part for understanding the product. When possible, links are provided to technical papers and Redbooks that supplement the broad coverage of the release documentation with in-depth examinations of particular product areas.

View links to additional information about:

- Programming instructions and examples
- Programming specifications

### Programming instructions and examples

- Java internationalization tutorial

An online tutorial that explains how to use the Java 2 SDK Internationalization API.

### Programming specifications

- Java 2 SDK, Standard Edition Documentation: Internationalization

The Java internationalization documentation from Sun Microsystems, including a list of supported locales and encodings.

- Java Specification Request 150, Internationalization Service for J2EE

The specification of the J2EE internationalization service that is currently being developed through the Java Community Process.

- W3C, Web Services Internationalization Task Force

The task force of the W3C's Internationalization Working Group responsible for investigating the internationalization of Web services, in particular, the dependence of Web services on language, culture, region, and locale-related contexts.

- Making the WWW truly World Wide

The W3C's effort to make World Wide Web technology work with the many writing systems, languages, and cultural conventions of the global community:

- developerWorks - Unicode

Articles on various subjects relating to Unicode, from IBM's developerWorks.

---

## Chapter 16. Application profiling

Application profiling enables you to configure multiple access intent policies on the same entity bean. Application profiling reflects the fact that different invocations against the same entity can require different kinds of support from the server run-time environment. For more information, see Application Profiling: Overview.

1. Assembling applications for application profiles. This topic describes how to configure tasks, create application profiles, and configure tasks on profiles.
2. Managing application profiles. This topic describes how to add and remove tasks from application profiles using the administrative console.
3. Using the TaskNameManager API. This topic describes how to programmatically set the current task name, but you should use this technique sparingly. Wherever possible, use the declarative method instead, which results in more portable function.

---

### Application profiling: Overview

Application profiling enables you to identify particular units of work to the WebSphere Application Server run time environment. The run time can tailor its support to the exact requirements of that unit of work. Access intent is currently the only run time component that makes use of the application profiling functionality. For example, you can configure one transaction to load an entity bean with strong update locks and configure another transaction to load the same entity bean without locks.

Application profiling introduces two new concepts in order to achieve this function: *tasks* and *profiles*.

**Tasks** A task is a named unit of work within a distributed application. *Unit of work* in this case means a unique path within the application that may or may not correspond to a transaction or activity session. The name of the path is typically assigned declaratively to the method of an enterprise bean, to a J2EE client or servlet. This point of configuration marks the head of a graph or subgraph identified by the name of the task; the task name flows from the head of the graph downstream on all subsequent IOP requests, identifying each subsequent invocation along the graph as belonging to the developer-configured task.

#### Profiles

A profile is simply a mapping of a task to a set of access intent policies that are configured on entity beans. When an invocation on a bean (whether by a finder method, a CMR getter, or a dynamic query) requires data to be retrieved from the back end system, the current task associated with the request is used to determine the exact requirement of the transaction. The same bean loads and behaves differently in the context of the task-to-profile mapping. Each profile provides the developer an opportunity to reconfigure the application's access intent. If a request is operating in the absence of a task, the run time environment uses either a method-level access intent (if any) or a bean-level default access intent.

### Application profiles

An application profile is the set of access intent policies that should be selectively applied for a particular unit of work (a transaction or *ActivitySession*).

Application profiling enables applications to run under different sets of policies depending on the active task under which the application is operating.

Consider an application that centralizes the student records for a school district. These records are frequently accessed by the school district's central office in order to generate reports. The report generation process would be optimized if it held no locks with the back end system, and if the records could be read into memory with as few back end operations as possible. Occasionally, however, the records are updated by the students' instructors. Without the ability to distinguish between transactions, the developer is forced to assume a worst-case scenario and, wishing to use pessimistic concurrency, lock the records for all transactions.

Using the application profiling service, the developer can configure in as many ways as necessary the access intent under which the students' records are loaded. Under one profile, the records can be configured with an exclusive pessimistic update intent, not only locking-out competing transactions but ensuring that the student is not removed from the system before the transaction completes. Under another profile, the records can be configured with an optimistic intent as part of an object graph that is read from the back end system in a single database operation. The task represented by the pessimistic profile receives the strong-locking semantics required for certain transactions, while the task represented by the optimistic profile receives the performance benefits appropriate for other transactions.

## Application profiling performance considerations

Application profiling enables assembly configuration techniques that improve your application runtime, performance and scalability. You can configure tasks that identify incoming requests, identify access intents determining concurrency and other data access characteristics, and profiles that map the tasks to the access intents. The capability to configure the application server can improve performance, efficiency and scalability, while reducing development and maintenance costs. The application profiling service has no tuning parameters, other than a checkbox for disabling the service if the service is not necessary. However, the overhead for the application profile service is small and should not be disabled, or unpredictable results can occur.

Access intents enable you to specify data access characteristics. The WebSphere run-time environment uses these hints to optimize the access to the data, by setting the appropriate isolation level and concurrency. Various access intent hints can be grouped together in an access intent policy.

In WebSphere Application Server, access intent policies are configured on methods; when an invocation of a method drives the activation of a bean, the access intent configured on that method controls the data access characteristics, such as concurrency. In WebSphere Business Integration Server Foundation, application profiling enables you to configure multiple access intent policies on the entity bean. WebSphere Business Integration Server Foundation chooses the appropriate access intent for the entity bean, using the task associated with the current unit of work (UOW). Some callers can load a bean with the intent to read data, while others can load the bean for update. The capability to configure the application server can improve performance, efficiency, and scalability, while reducing development and maintenance costs.

Access intents enable the EJB container to be configured providing optimal performance based on the specific type of enterprise bean used. Various access intent hints can be specified declaratively at deployment time to indicate to WebSphere resources, such as the container and persistence manager, to provide the appropriate access intent services for every EJB request.

The application profiling service improves overall entity bean performance and throughput by fine tuning the runtime behavior. The application profiling service enables EJB optimizations to be customized for multiple user access patterns without resorting to "worst case" choices, such as pessimistic update on a bean accessed with the `findByPrimaryKey` method, regardless of whether the client needs it for read or for an update.

Application profiling provides the capability to define the following hierarchy: **Container Tasks** > **Application Profiles** > **Access Intent Policies** > **Access Intent Overrides**. Container tasks identify UOW and are associated with a method or a set of methods. When a method associated with the task is invoked, the task name is propagated with the request. For example, a UOW refers to a unique path within the application that can correspond to a transaction or activity session. The name of the task is assigned declaratively to a J2EE client or servlet, or to the method of an enterprise bean. The task name identifies the starting point of a call graph or subgraph; the task name flows from the starting point of the graph downstream on all subsequent IOP requests, identifying each subsequent invocation along the graph as belonging to the configured task. As a best practice, wherever a UOW starts, for example, a transaction or an `ActivitySession`, assign a task to that starting point and make sure the task remains associated with that UOW. It is best to configure the application server so that no more than one transaction is associated with more than one task.

The application profile service associates the propagated tasks with access intent policies. When a bean is loaded and data is retrieved, the characteristics used for the retrieval of the data are dictated by the application profile. The application profile configures the access intent policy and the overrides that should be used to access data for a specific task.

Access intent policies determine how beans are loaded for specific tasks and how data is accessed during the transaction. The access intent policy is a named group of access intent hints. The hints can be used, depending on the characteristics of the database and resource manager. Various access intent hints applied to the data access operation govern data integrity. The general rule is, the more data integrity, the more overhead. More overhead causes lower throughput and the opportunity for simultaneous data access from multiple clients.

If specified, access intent overrides provide further configuration for the access intent policy.

## Best practices

Application profiling is effective in a variety of different scenarios. The following are example situations where application profiling is useful

- **The same bean is loaded with different data access patterns**

The same bean or set of beans can be reused across applications, but each of those applications has differing requirements for the bean or for beans within the invocation graph. One application can require that beans be loaded for update, while another application requires beans be loaded for read only. Application profiling enables deploy time configuration for beans to distinguish between EJB loading requirements.

- **Different clients have different data access requirements**

The same bean or set of beans can be used for different types of client requests. When those clients have different requirements for the bean, or for beans within the invocation graph, application profiling can be used to tailor the bean loading characteristics to the requirements of the client. One client can require beans be loaded for update, while another client requires beans be loaded for read only. Application profiling enables deploy time configuration for beans to distinguish between EJB loading requirements.

## Monitoring tools

You can use the Tivoli Performance Viewer, database and logs as monitoring tools.

You can use the Tivoli Performance Viewer to monitor various metrics associated with beans in an application profiling configuration. The following sections describe at a high level the Tivoli Performance Viewer metrics that reflect changes when access intents and application profiling are used:

- **Collection scope**

The enterprise beans group contains EJB life cycle information, either a cumulative value for a group of beans, or for specific beans. You can monitor this information to determine the difference between using the `ActivitySession` scope versus the transaction scope. For the transaction scope, depending on how the container transactions are defined, `activates` and `passivates` can be associated with method invocations. The application could use the `ActivitySession` scope to reduce the frequency of `activates` and `passivates`. For more information, see "Using the `ActivitySession` service."

- **Collection increment**

The enterprise beans group contains EJB life cycle information, either a cumulative value for a group of beans, or for specific beans. You can monitor `Num Activates` to watch the number of enterprise beans activated for a particular `findByPrimaryKey` operation. For example, if the collection increment is set to 10, rather than the default 25, the `Num Activates` value shows 25 for the initial `findByPrimaryKey`, before any result set iterator runs. If the number of `activates` rarely exceeds the collection increment, consider reducing the collection increment setting.

- **Resource manager prefetch increment**

The resource manager prefetch increment is a hint acted upon by the database engine to depend upon the database. The Tivoli Performance Viewer does not have a metric available to show the effect of the resource manager prefetch increment setting.

- **Read ahead hint**

The enterprise beans group contains EJB life cycle information, either a cumulative value for a group of beans, or for specific beans. You can monitor *Num Activates* to watch the number of enterprise beans activated for a particular request. If a read ahead association is not in use, the *Num Activates* value shows a lower initial number. If a read ahead association is in use, the *Num Activates* value represents the number of activates for the entire call graph.

**Database tools** are helpful in monitoring the different bean loading characteristics that introduce contention and concurrency issues. These issues can be solved by application profiling, or can be made worse by the misapplication of access intent policies.

Database tools are useful for monitoring locking and contention characteristics, such as locks, deadlocks and connections open. For example, for locks the DB2 Snapshot Monitor can show statistics for lock waits, lock time-outs and lock escalations. If excessive lock waits and time-outs are occurring, application profiling can define specific client tasks that require a more string level of locking, and other client tasks that do not require locking. Or, a different access intent policy with less restrictive locking could be applied. After applying this configuration change, the snapshot monitor shows less locking behavior. Refer to information about the database you are using on how to monitor for locking and contention.

The **application server logs** can be monitored for information about rollbacks, deadlocks, and other data access or transaction characteristics that can degrade performance or cause the application to fail.

---

## Application profiling tasks

Tasks are named units of work. They are the mechanism by which the run time environment determines which access intent policies to apply when an entity bean's data is loaded from the back end system.

Application profiles enable developers to configure an entity bean with multiple access intent policies; if there are  $n$  instances of profiles in a given application, each bean can be configured with as many as  $n$  access intent policies.

A task is a unit of work that is given a name by a developer. A task is assigned to any thread executing within a J2EE component, then propagated implicitly across all IOP requests. The WebSphere Application Server run time environment queries the task at the invocation of any entity bean, and establishes the appropriate access intent policy with which an entity instance is associated. A task typically corresponds to the execution of a concrete and high-level job within the application.

If an entity bean is loaded in a unit of work that is not associated with a task, or is associated with a task that is unassociated with an application profile, the default bean-level access intent or the method-level access intent configuration is applied. If a unit of work is associated with a task that is configured with an application profile, the bean-level access intent configuration within the appropriate application profile is applied.

For example, consider a school district application that calls through a session bean in order to interact with student records. One method on the session bean allows administrators to modify the students' records; another method supports student requests to view their own records. Without application profiling, the two tasks would operate anonymously and the run time environment would be unable to distinguish work operating on behalf of one task or the other. To optimize the application, a developer can configure one of the methods on the session bean with the task "updateRecords" and the other method on the session bean with the task "readRecords". When registered with an application profile that has the student bean configured with the appropriate locking access intent, the "updateRecords" task is assured that it is not unnecessarily blocking transactions that need to only read the records.

Tasks can be configured to be managed by the container or to be programmatically established by the application. Container managed tasks can be configured on servlets, Java Server Pages (JSPs), application clients, and the methods of Enterprise Java Beans (EJBs). Application managed tasks are configured on all J2EE components.

---

## Assembling applications for application profiling

Application profiling enables multiple access intent policies to be configured on the same entity bean, each specified for a particular unit of work. You can use the one of the default policies or create your own, as described in the topic, *Creating a custom access intent policy*.

1. Configuring tasks. Declaratively configure tasks as described in the following topics:
  - Configuring container-managed tasks for Enterprise Java Beans.
  - Configuring container-managed tasks for web components.
  - Configuring container-managed tasks for application clients.

On rare occasions, you might find it necessary to configure tasks *programmatically*. Application profiling supports this requirement with a simple interface that enables both overriding of the current task associated with the thread of execution, and resetting of the current task to the original task. See *Using the TaskNameManager interface*.

2. Creating an application profile.

## Automatic configuration of application profiling

Application profiling requires accurate knowledge of an application's transactional configuration and the interaction of the application with its persistent state during the course of each transaction.

The WebSphere Integration Server assembly tooling includes a static analysis engine that can assist you in configuring application profiling. The tool examines the compiled classes and the deployment descriptor of a Java 2 Enterprise Edition (J2EE) application to determine the entry point of transactions, calculate the set of entities enlisted in each transaction, and determine whether the entities are read or updated during the course of each identified transaction.

You can execute the analysis in either *closed world* or *open world* mode. A closed-world analysis assumes that all possible clients of the application are included in the analysis and that the resulting analysis is complete and correct. The results of a closed-world analysis report the set of all transactions that can be invoked by a web, JMS, or application client. The results exclude many potential transactions that never execute at run time.

An open-world analysis assumes that not all clients are available for analysis or that the analysis cannot return complete or accurate results. An open-world analysis returns the complete set of possible transactions.

The results of an analysis persist as an application profiling configuration. The tool establishes container managed tasks for servlets, Java Server Pages (JSPs), application clients, and Message Driven Beans (MDBs). Application profiles for the tasks are constructed with the appropriate access intent for the entities enlisted in the transaction represented by the task. However, in practice, there are many situations where the tool returns at best incomplete results. Not all applications are amenable to static analysis. Some factory and command patterns make it impossible to determine the call graphs. The tool does not support the analysis of *ActivitySessions*.

You should examine the results of the analysis very carefully. In many cases you must manually modify them to meet the requirements of the application. However, the tool can be an effective starting place for most applications and may offer a complete and quick configuration of application profiles for some applications.



## Automatically configuring application profiles and tasks

You can automatically configure application profiling for an application through static analysis.

1. Start the Assembly Toolkit.
2. Optional: Open the J2EE perspective to work with J2EE projects. Click **Window > Open Perspective > Other > J2EE**.
3. Optional: Open the J2EE Hierarchy view. Click **Window > Show View > J2EE Hierarchy**. Other helpful views include the Project Navigator view (**Window > Show View > Other > J2EE > Project Navigator**) and the Navigator view (**Window > Show View > Navigator**).
4. Create a new application EAR file or edit an existing one.

**Note:** Ensure that you have Enabled Server Target support for J2EE projects and that you set the target server as *Integration Server v5.1* or *WebSphere Application Server EE v5.0*.

For example, to change attributes of an existing application, use the import wizard to import an EAR file. To start the import wizard:

- a. Select **File > Import > EAR file > Next**
  - b. Select the EAR file.
  - c. In the *Target server* field, select *Integration Server v5.1* or *WebSphere Application Server EE v5.0*
  - d. Select **Finish**
5. Be sure that the application and its modules successfully compile. To include Java Server Pages (JSPs) in the analysis, you must precompile the JSPs. Also, be sure that you have configured all transactional attributes before analyzing.
  6. In the J2EE Hierarchy view of the J2EE perspective, right-click the **Enterprise Application** to configure, then select **Open With > Deployment Descriptor Editor**. A property dialog notebook for the application project is displayed in the property pane.
  7. In the property pane, select the **Extended Services** tab.
  8. Beneath the *Application Profiles* table, select **Auto...**
  9. Select the projects to be analyzed and configured. Select **Next**.
  10. To limit the returned results of the analysis, choose *closed world analysis*. Closed world analysis generates application profiles only if a client entry point in a message driven bean (MDB), servlet, JSP, or application client is resolved that begins a transaction and enlists entities. If closed world is not selected, the analysis returns the set of application profiles for all possible transactions represented by the application.

**Note:** At this point you can also choose the **Clean** attribute. If you set this attribute, the existing configuration of selected modules is removed and the new configuration is applied fresh. If you do not select this option, the new configuration is merged into the existing configuration.

11. Choose the concurrency for the default configuration of the access intent for generated application profiles.
12. Select **Analyze > Next**.
13. Examine the results of the analysis. Each top-level entry in the table represents a transaction identified by the analysis. The nested entries represent the callers of the transaction, the entities enlisted by the transaction, and the attributes read or modified during the course of the transaction.
14. Select **Finish** to automatically configure the container-managed tasks and application profiles represented by the analysis.

## Applying profile-scoped access intent policies to entity beans

Configure entities with access intent for an application profile

1. Start the Assembly Toolkit.



2. Optional: Open the J2EE perspective to work with J2EE projects. Click **Window > Open Perspective > Other > J2EE**.
3. Optional: Open the J2EE Hierarchy view. Click **Window > Show View > J2EE Hierarchy**. Other helpful views include the Project Navigator view (**Window > Show View > Other > J2EE > Project Navigator**) and the Navigator view (**Window > Show View > Navigator**).
4. Create a new application EAR file or edit an existing one.

**Note:** Ensure that you have Enabled Server Target support for J2EE projects and that you set the target server as *Integration Server v5.1* or *WebSphere Application Server EE v5.0*.

For example, to change attributes of an existing application, use the import wizard to import an EAR file. To start the import wizard:

- a. Select **File > Import > EAR file > Next**
  - b. Select the EAR file.
  - c. In the *Target server* field, select *Integration Server v5.1* or *WebSphere Application Server EE v5.0*
  - d. Select **Finish**
5. In the J2EE Hierarchy view of the J2EE perspective, right-click the **EJB module** for the application profile instance, then select **Open With > Deployment Descriptor Editor**. A property dialog notebook for the EJB project is displayed in the property pane.
  6. Select the **Extended Access** tab.
  7. Select the **application profile** for which you want to specify the access intent.
  8. Beneath the **Access Intent for Entities 2.x (Profile Level)** panel, select **Add...**
  9. Select the **entities** to configure and click **Next...**
  10. Select the **access intent policy** to apply. Select **Read Ahead Hint** if a read ahead hint is desired.
  11. Select **Next**.
  12. Optional: Specify the **collection scope**

#### **Transaction**

This is the default. Collections of entities cannot be used beyond the scope of the transaction in which you create the collection.

#### **ActivitySession**

Collections of entities cannot be used beyond the scope of the *ActivitySession* in which you create the collection. The collection can be used in a new transaction if that transaction is nested under the original *ActivitySession*, although you might have to reload the object by querying the underlying data store.

13. Optional: Specify the **collection increment**. Specify a valid integer to define the chunks that populate a remote collection. This value only applies to *remote* collections and is ignored by local collections. The default for access types that result in U locks is 1. Otherwise, the default is 25.
14. Optional:
15. Specify the **resource manager prefetch increment**. Specify a valid integer to set as the fetch size on the JDBC statement when you execute queries for a bean type. The default is 0.
16. Select **Next**.
17. If you selected *read ahead*, choose the **preload path**.
18. Select **Finish** to apply.
19. Select **OK**.

## **Creating a custom access intent policy**

Define a custom access intent policy which can be configured for 2.0 entity beans.

1. Start the Assembly Toolkit.

2. Optional: Open the J2EE perspective to work with J2EE projects. Click **Window > Open Perspective > Other > J2EE**.
3. Optional: Open the J2EE Hierarchy view. Click **Window > Show View > J2EE Hierarchy**. Other helpful views include the Project Navigator view (**Window > Show View > Other > J2EE > Project Navigator**) and the Navigator view (**Window > Show View > Navigator**).
4. Create a new application EAR file or edit an existing one.

**Note:** Ensure that you have Enabled Server Target support for J2EE projects and that you set the target server as *Integration Server v5.1* or *WebSphere Application Server EE v5.0*.

For example, to change attributes of an existing application, use the import wizard to import an EAR file. To start the import wizard:

- a. Select **File > Import > EAR file > Next**
  - b. Select the EAR file.
  - c. In the *Target server* field, select *Integration Server v5.1* or *WebSphere Application Server EE v5.0*
  - d. Select **Finish**
5. In the J2EE Hierarchy view of the J2EE perspective, right-click the **EJB module**, then select **Open With > Deployment Descriptor Editor**. A property dialog notebook for the Enterprise Java Bean (EJB) project is displayed in the property pane.
  6. In the property pane, select the **Extended Access** tab.
  7. Beneath the **Defined Access Intent Policies** panel, select **Add**.
  8. Specify a **unique name** by which the policy is referenced when applied to entity beans.
  9. Optional: Specify a **description** of the policy.
  10. Specify an **access type**.
  11. Specify the **collection scope**.

#### **Transaction**

This is the default. Collections of entities cannot be used beyond the scope of the transaction in which the collection is created.

#### **ActivitySession**

Collections of entities cannot be used beyond the scope of the *ActivitySession* in which the collection is created. The collection can be used in a new transaction if that transaction is nested under the original *ActivitySession*, although you might need to reload the object by querying the underlying data store.

12. Specify the **collection increment**. Specify a valid integer to define the chunks that populate a remote collection. This value only applies to *remote* collections and is ignored by local collections. The default value for access types that result in U locks is 1. Otherwise, the default is 25.
13. Specify the **resource manager prefetch increment**. Specify a valid integer to set as the fetch size on the JDBC statement when executing queries for a bean type. The default value is 0.
- 14.

Applying access intent policies to entity beans.

## **Creating an application profile**

An application profile contains a set of access intent policies applied to an application's entity beans. The access intent policies are only applied for requests that are associated with tasks configured on the application profile.

1. Start the Assembly Toolkit.
2. Optional: Open the J2EE perspective to work with J2EE projects. Click **Window > Open Perspective > Other > J2EE**.

3. Optional: Open the J2EE Hierarchy view. Click **Window > Show View > J2EE Hierarchy**. Other helpful views include the Project Navigator view (**Window > Show View > Other > J2EE > Project Navigator**) and the Navigator view (**Window > Show View > Navigator**).
4. Create a new application EAR file or edit an existing one.

**Note:** Ensure that you have Enabled Server Target support for J2EE projects and that you set the target server as *Integration Server v5.1* or *WebSphere Application Server EE v5.0*.

For example, to change attributes of an existing application, use the import wizard to import an EAR file. To start the import wizard:

- a. Select **File > Import > EAR file > Next**
  - b. Select the EAR file.
  - c. In the *Target server* field, select *Integration Server v5.1* or *WebSphere Application Server EE v5.0*
  - d. Select **Finish**
5. In the J2EE Hierarchy view of the J2EE perspective, right-click the **EJB module**, then select **Open With > Deployment Descriptor Editor**. A property dialog notebook for the Enterprise Java Bean (EJB) project is displayed in the property pane.
  6. In the property pane, select the **Extended Access** tab.
  7. Beneath the **Application Profiles** table, select **Add...**
  8. Select or type the **name** of the task for which this profile applies. You cannot map the task to any other profile within the module. The task name should have been configured already as a task on a web component (container managed task, application managed task), an application client (container managed task, application managed task), or the method of an EJB (container managed task, application managed task).
  9. Optional: Specify a **description** of the task.
  10. Select **Next**.
  11. Select the entities that are enlisted in the unit of work represented by the application profile. You can add additional entities as a separate task after the profile has been created.
  12. Select **Finish**.

Configuring access intent for application profiles.

## Configuring container managed tasks for application clients

Configure an application client's container-managed task to associate requests from the client with an application profile.

1. Start the Assembly Toolkit.
2. Optional: Open the J2EE perspective to work with J2EE projects. Click **Window > Open Perspective > Other > J2EE**.
3. Optional: Open the J2EE Hierarchy view. Click **Window > Show View > J2EE Hierarchy**. Other helpful views include the Project Navigator view (**Window > Show View > Other > J2EE > Project Navigator**) and the Navigator view (**Window > Show View > Navigator**).
4. Create a new application EAR file or edit an existing one.

**Note:** Ensure that you have Enabled Server Target support for J2EE projects and that you set the target server as *Integration Server v5.1* or *WebSphere Application Server EE v5.0*.

For example, to change attributes of an existing application, use the import wizard to import an EAR file. To start the import wizard:

- a. Select **File > Import > EAR file > Next**
- b. Select the EAR file.
- c. In the *Target server* field, select *Integration Server v5.1* or *WebSphere Application Server EE v5.0*

- d. Select **Finish**
5. In the J2EE Hierarchy view of the J2EE perspective, right-click the **Application Client module**, then select **Open With > Deployment Descriptor Editor**. A property dialog notebook for the application client project is displayed in the property pane.
6. Select the **Extended Services** tab.
7. Enter the **name** and **description** of the task in the **Container-Managed Task** section. The task name is mapped to application profiles and used by the run time to determine the appropriate access intent to use for enlisted entities. Task names do not have to be unique within an application. However, task names should be shared consciously and conservatively. At run time, all tasks with the same name are treated the same way, regardless of where you configured the task.  
The description is provided for your convenience, it is not used by the run time environment.
8. Select **OK**.

## Configuring container managed tasks for web components

Configure a web component's container managed task to associate requests from a servlet or Java Server Page (JSP) with an application profile.

1. Start the Assembly Toolkit.
2. Optional: Open the J2EE perspective to work with J2EE projects. Click **Window > Open Perspective > Other > J2EE**.
3. Optional: Open the J2EE Hierarchy view. Click **Window > Show View > J2EE Hierarchy**. Other helpful views include the Project Navigator view (**Window > Show View > Other > J2EE > Project Navigator**) and the Navigator view (**Window > Show View > Navigator**).
4. Create a new application EAR file or edit an existing one.

**Note:** Ensure that you have Enabled Server Target support for J2EE projects and that you set the target server as *Integration Server v5.1* or *WebSphere Application Server EE v5.0*.

For example, to change attributes of an existing application, use the import wizard to import an EAR file. To start the import wizard:

- a. Select **File > Import > EAR file > Next**
- b. Select the EAR file.
- c. In the *Target server* field, select *Integration Server v5.1* or *WebSphere Application Server EE v5.0*
- d. Select **Finish**
5. In the J2EE Hierarchy view of the J2EE perspective, right-click the **WEB module** for the web component instance, then select **Open With > Deployment Descriptor Editor**. A property dialog notebook for the web project is displayed in the property pane.
6. Select the **Extended Services** tab.
7. Select the servlet or JSP that you want to change.
8. Enter the **name** and **description** of the task in the **Container-Managed Task** section. The task name is mapped to application profiles and used by the run time to determine the appropriate access intent to use for enlisted entities. Task names do not have to be unique within an application. However, task names should be shared consciously and conservatively. At run time, all tasks with the same name are treated the same way, regardless of where you configured the task.  
The description is provided for your convenience, it is not used by the run time environment.
9. Select **OK**.

## Configuring container managed tasks for Enterprise Java Beans

Configure an Enterprise Java Bean's (EJB) container managed tasks to associate requests from the bean with application profiles.

1. Start the Assembly Toolkit.
2. Optional: Open the J2EE perspective to work with J2EE projects. Click **Window > Open Perspective > Other > J2EE**.
3. Optional: Open the J2EE Hierarchy view. Click **Window > Show View > J2EE Hierarchy**. Other helpful views include the Project Navigator view (**Window > Show View > Other > J2EE > Project Navigator**) and the Navigator view (**Window > Show View > Navigator**).
4. Create a new application EAR file or edit an existing one.

**Note:** Ensure that you have Enabled Server Target support for J2EE projects and that you set the target server as *Integration Server v5.1* or *WebSphere Application Server EE v5.0*.

For example, to change attributes of an existing application, use the import wizard to import an EAR file. To start the import wizard:

- a. Select **File > Import > EAR file > Next**
  - b. Select the EAR file.
  - c. In the *Target server* field, select *Integration Server v5.1* or *WebSphere Application Server EE v5.0*
  - d. Select **Finish**
5. In the J2EE Hierarchy view of the J2EE perspective, right-click the **EJB module** for the bean instance, then select **Open With > Deployment Descriptor Editor**. A property dialog notebook for the EJB project is displayed in the property pane.
  6. Select the **Extended Access** tab.
  7. Beneath the **Container-Managed Tasks** table, select **Add...**
  8. Select the **bean** for which you want to configure the task.
  9. Select **Next**.
  10. Select the **method** for which you want to configure the task.
  11. Select **Next**.
  12. Enter the **name** and **description** of the task. The task name is mapped to application profiles and used by the run time to determine the appropriate access intent to use for enlisted entities. Task names do not have to be unique within an application. However, task names should be shared consciously and conservatively. At run time, all tasks with the same name are treated the same way, regardless of where you configured the task.  
The description is provided for your convenience, it is not used by the run time environment.
  13. Select **OK**.

## Configuring application managed tasks for application clients

Configure application managed tasks for an application client to associate requests from the client with application profiles.

1. Start the Assembly Toolkit.
2. Optional: Open the J2EE perspective to work with J2EE projects. Click **Window > Open Perspective > Other > J2EE**.
3. Optional: Open the J2EE Hierarchy view. Click **Window > Show View > J2EE Hierarchy**. Other helpful views include the Project Navigator view (**Window > Show View > Other > J2EE > Project Navigator**) and the Navigator view (**Window > Show View > Navigator**).
4. Create a new application EAR file or edit an existing one.

**Note:** Ensure that you have Enabled Server Target support for J2EE projects and that you set the target server as *Integration Server v5.1* or *WebSphere Application Server EE v5.0*.

For example, to change attributes of an existing application, use the import wizard to import an EAR file. To start the import wizard:

- a. Select **File > Import > EAR file > Next**

- b. Select the EAR file.
  - c. In the *Target server* field, select *Integration Server v5.1* or *WebSphere Application Server EE v5.0*
  - d. Select **Finish**
5. In the J2EE Hierarchy view of the J2EE perspective, right-click the **Client module**, then select **Open With > Deployment Descriptor Editor**. A property dialog notebook for the application client project is displayed in the property pane.
  6. In the property pane, select the **Extended Services** tab.
  7. Beneath the **Application-Managed Tasks** panel, select **Add...**
  8. Specify the name of the **task reference**. The task reference name is used programmatically by the application client. The task reference is the logical representation of the task that is used by the run time environment.
  9. Enter the **name** and **description** of the task. The task name is mapped to application profiles and used by the run time to determine the appropriate access intent to use for enlisted entities. Task names do not have to be unique within an application. However, task names should be shared consciously and conservatively. At run time, all tasks with the same name are treated the same way, regardless of where you configured the task.  
The description is provided for your convenience, it is not used by the run time environment.
  10. Select **OK**.

## Configuring application-managed tasks for web components

Configure web components' application managed tasks to associate requests from a servlet or Java Server Page (JSP) with application profiles.

1. Start the Assembly Toolkit.
2. Optional: Open the J2EE perspective to work with J2EE projects. Click **Window > Open Perspective > Other > J2EE**.
3. Optional: Open the J2EE Hierarchy view. Click **Window > Show View > J2EE Hierarchy**. Other helpful views include the Project Navigator view (**Window > Show View > Other > J2EE > Project Navigator**) and the Navigator view (**Window > Show View > Navigator**).
4. Create a new application EAR file or edit an existing one.

**Note:** Ensure that you have Enabled Server Target support for J2EE projects and that you set the target server as *Integration Server v5.1* or *WebSphere Application Server EE v5.0*.

For example, to change attributes of an existing application, use the import wizard to import an EAR file. To start the import wizard:

- a. Select **File > Import > EAR file > Next**
  - b. Select the EAR file.
  - c. In the *Target server* field, select *Integration Server v5.1* or *WebSphere Application Server EE v5.0*
  - d. Select **Finish**
5. In the J2EE Hierarchy view of the J2EE perspective, right-click the **Web module** for the web component instance, then select **Open With > Deployment Descriptor Editor**. A property dialog notebook for the web project is displayed in the property pane.
  6. In the property pane, select the **Extended Services** tab.
  7. Select the servlet or JSP that you want to change.
  8. Beneath the **Application-Managed Tasks** table, select **Add...**
  9. Specify the name of the **task reference**. The task reference name is used programmatically by the servlet or JSP. The task reference is the logical representation of the task that is used by the run time environment.
  10. Enter the **name** and **description** of the task. The task name is mapped to application profiles and used by the run time to determine the appropriate access intent to use for enlisted entities. Task



names do not have to be unique within an application. However, task names should be shared consciously and conservatively. At run time, all tasks with the same name are treated the same way, regardless of where you configured the task.

The description is provided for your convenience, it is not used by the run time environment.

11. Select **OK**.

## Configuring application managed tasks for Enterprise Java Beans

Configure an Enterprise Java Bean's (EJB) application managed tasks to associate requests from the bean with application profiles.

1. Start the Assembly Toolkit.
2. Optional: Open the J2EE perspective to work with J2EE projects. Click **Window > Open Perspective > Other > J2EE**.
3. Optional: Open the J2EE Hierarchy view. Click **Window > Show View > J2EE Hierarchy**. Other helpful views include the Project Navigator view (**Window > Show View > Other > J2EE > Project Navigator**) and the Navigator view (**Window > Show View > Navigator**).
4. Create a new application EAR file or edit an existing one.

**Note:** Ensure that you have Enabled Server Target support for J2EE projects and that you set the target server as *Integration Server v5.1* or *WebSphere Application Server EE v5.0*.

For example, to change attributes of an existing application, use the import wizard to import an EAR file. To start the import wizard:

- a. Select **File > Import > EAR file > Next**
  - b. Select the EAR file.
  - c. In the *Target server* field, select *Integration Server v5.1* or *WebSphere Application Server EE v5.0*
  - d. Select **Finish**
5. In the J2EE Hierarchy view of the J2EE perspective, right-click the **EJB module** for the bean instance, then select **Open With > Deployment Descriptor Editor**. A property dialog notebook for the EJB project is displayed in the property pane.
  6. Select the **Extended Access** tab.
  7. Beneath the **Application-Managed Tasks** table, select **Add...**
  8. Select the bean for which you want to configure the task.
  9. Type the **name** of the task reference. The task reference name is used programmatically by the bean. The task reference is the logical representation of the task used by the run time environment.
  10. Enter the **name** and **description** of the task. The task name is mapped to application profiles and used by the run time to determine the appropriate access intent to use for enlisted entities. Task names do not have to be unique within an application. However, task names should be shared consciously and conservatively. At run time, all tasks with the same name are treated the same way, regardless of where you configured the task.

The description is provided for your convenience, it is not used by the run time environment.

11. Select **OK**.

## Specifying target servers for J2EE projects

The J2EE tooling of the Assembly Toolkit and the WebSphere Studio Application Developer Integration Edition enables you to target a specific server during your application development. For example, if you want to take advantage of the features of JDK 1.4.1, which is used as the run time environment for WebSphere Application Server Version 5.1, your applications require different class path entries than those that were used previously. With this server targeting feature, you can now easily update your class path



entries appropriately for running on WebSphere Application Server V5.1 using the JDK 1.4 run time environment. But more than just supporting JDK 1.4.1, targeting a server is an easy way to specify the JAR files available at run time for a specific server.

When a target server is specified for a project, a *.server* file is added to the project resources indicating the target server.

After you enable the server targeting feature in the preferences, the J2EE project creation and import wizards include an additional option where you can specify a target server for the project. The list of target servers that you can choose from is filtered based on the J2EE level of the application as well as the J2EE module type. For example, for an EJB project, only servers that support Enterprise Java Beans are displayed.

For more information about target servers, refer to the WebSphere Studio Application Developer Integration Edition documentation under the following topics:

- Specifying target servers for J2EE projects
- Modifying the target server for J2EE projects
- Synchronizing the target servers for J2EE projects

To enable the server targeting support:

1. Select **Window > Preferences** . The **Preferences** window opens.
2. On the **J2EE** page of the preferences, select the **Enable** radio button in the **Server Targeting Support** section.
3. Select **OK**.

## Modifying the target server for J2EE projects

If you enabled the server targeting feature in the J2EE preferences, you can switch the specified target server for existing EAR projects or EAR module projects. When a target server is specified for a project, a *.server* file is added to the project resources indicating the target server.

For more information about target servers, refer to the WebSphere Studio Application Developer Integration Edition documentation under the following topics:

- Specifying target servers for J2EE projects
  - Modifying the target server for J2EE projects
  - Synchronizing the target servers for J2EE projects
1. In the J2EE Hierarchy view of the J2EE perspective, right-click the enterprise application or module project, and select Target Server > Modify from the pop-up menu. The Modify Target Server dialog box opens.
  2. In the Target Server field, select the server that you want to target. If your project needs to be supported by earlier versions of this product, select the No target server specified option.
  3. Select the Update EAR modules and utility projects to the selected target server to synchronize the target server between the EAR and all modules. This ensures that the enterprise application project and all module and utility projects, as well as Web application library projects, all have the same target server.
- Click OK.

---

## Managing application profiles

Manage your application profiles using the administrative console. From the console, you can add tasks to, and remove tasks from, application profiles.

1. Start the administrative console.

2. Select **Applications > Applications > *application\_name* > Application Profile > *profile\_name* > Tasks**.
3. On the Tasks collection page, you can add new tasks to the profile, delete tasks, edit current task settings, and so on.  
Note that no task can, within the scope of an application, be configured on more than one application profile. In such a situation, your application cannot be restarted until you correct the configuration.
4. Save your configuration.
5. Restart the application in order for your changes to take effect.

## Using the TaskNameManager interface

You can declaratively configure container managed tasks for J2EE web components, application clients, and Enterprise Java Beans (EJBs). On rare occasions, you might find it necessary to *programmatically* set the current task name. Application profiling supports this requirement with a simple interface that enables both overriding of the current task associated with the thread of execution, and resetting of the current task with the original task.

Application profiling does not support queries of the task that is in operation at run time. Instead, applications interact with logical task names that are declaratively configured as application managed tasks. Logical references enable the actual task name to be changed without having to recompile applications.

Wherever possible, avoid setting tasks programmatically. The declarative method results in more portable function that can be easily adjusted without requiring redevelopment and recompilation.

1. Configure application-managed tasks. Application profiling requires that a task name reference be declared for any task that is to be set programmatically. Task name references introduce a level of indirection so that the actual task set at run time can be adjusted by reassembly without requiring recoding or recompilation. Any attempt to set a task name that is undeclared as a task reference results in the raising of an exception.

Configure application-managed tasks as described in the following topics:

- Configuring application managed tasks for web components.
- Configuring application managed tasks for application clients.
- Configuring application managed tasks for EJBs.

2. Perform a JNDI lookup on the TaskNameManager interface:

```
InitialContext ic = new InitialContext();
TaskNameManager tnManager = ic.lookup
("java:comp/websphere/AppProfile/TaskNameManager");
```

The *TaskNameManager* interface is not bound into the namespace if the application profiling service is disabled.

3. Set the task name:

```
try {
tnManager.setTaskName("updateAccount");
}
catch (IllegalTaskNameException e) {
// task name reference not configured. Handle error.
}
// . . .
rnManager.resetTaskName();
```

Resetting the task name undoes the effects of any *setTaskName()* method operations and reestablishes whatever task name was current when the component began execution. If the *setTaskName()* method has not been called, the *resetTaskName()* method has no effect.

## TaskNameManager interface

The TaskNameManager interface is available to all J2EE components using the following JNDI lookup:

```
java:comp/websphere/AppProfile/TaskNameManager
```

```
package com.ibm.websphere.appprofile;
```

```
/**
 * The TaskNameManager is the programmatic interface
 * to the application profiling function. Using this interface,
 * programmers can set the current task name on the
 * thread of execution. The task name must have been
 * configured in the deployment descriptors as a task
 * reference associated with a task. The set task
 * name's scope is the duration of the method
 * invocation in the EJB and Web components and for
 * the duration of the client process, or until the
 * resetTaskName() method is invoked.
 */
public interface TaskNameManager {

/**
 * Set the thread's current task name to the specified
 * parameter. The task name must have been configured as
 * a task reference with a corresponding task or the
 * IllegalArgumentException exception is thrown.
 */
public void setTaskName(String taskName) throws IllegalArgumentException;

/**
 * Sets the thread's task name to the value that was set
 * at, or imported into, the beginning of the method
 * invocation (for EJB and Web components) or process
 * (for J2EE clients).
 */
public void resetTaskName();

}
```

## Application profiling exceptions

The following exceptions are thrown in response to various illegal actions related to application profiling:

### **com.ibm.ws.exception.RuntimeWarning**

This exception is thrown when the application is started, if the application is configured incorrectly.

The startup is consequently terminated. Some examples of misconfiguration include:

- A task configured on two different application profiles.
- A method configured with two different task run-as policies .

### **com.ibm.websphere.appprofile.IllegalTaskNameException**

This exception is raised if an application attempts to programmatically set a task when that task has not been configured as a task name reference.

## Application profiling service settings

Use this page to enable or disable the application profiling service.

Applications that are configured to use the application profiling service will not start successfully unless the application profiling service is enabled.

To view this administrative console page, click **Servers > Application Servers > server\_name > Application Profiling Service**.

### **Startup**

Specifies whether the server will attempt to start the application profiling service.

**Default  
Range**

Selected  
**Selected**

When the application server starts, it attempts to start the application profiling service automatically.

**Cleared**

This option is unavailable. The application profiling service cannot be disabled.

## Application profile collection

Use this page to manage application profiles.

An application profile is a set of policies that are to be applied during the execution of an enterprise bean and a set of tasks that are associated with that profile. Mapping tasks to application profiles will control which access intent policies are applied at run time for the units of work that correspond to a particular task.

To view this administrative console page, click **Applications > Applications > *application\_name* > Application Profile**.

### Name

The name of the application profile.

The name must be unique; multiple profiles cannot share the same name.

**Data type** String

### Description

A description of the application profile.

**Data type** String

## Application profile settings

Use this page to modify application profile settings.

To view this administrative console page, click **Applications > Applications > *application\_name* > Application Profile > *application\_profile\_name***.

### **Name:**

The name of the application profile.

The name must be unique; multiple profiles cannot share the same name.

**Data type** String

### **Description:**

A description of the application profile.

**Data type** String

### **Task collection:**

Use this page to manage tasks.

Requests associated with any of the configured tasks operate under the access-intent policies that are configured with the profile. A task can be configured on only *one* application profile.

To view this administrative console page, click **Applications > Applications > application\_name > Application Profile > application\_profile\_name > Tasks**.

*Name:*

The name of the task.

The task name must be unique among the set of application profiles.

**Data type** String

*Description:*

A description of the task.

**Data type** String

*Task settings:*

Use this page to modify task settings.

To view this administrative console page, click **Applications > Applications > application\_name > Application Profile > application\_profile\_name > Tasks > task\_name**.

*Name:*

The name of the task.

The task name must be unique among the set of application profiles.

**Data type** String

*Description:*

A description of the task.

**Data type** String

---

## Chapter 17. Using Business Rule Beans

This topic provides a brief overview of the steps involved in externalizing Business Rule Beans. To gain an understanding of business rules and Business Rule Beans (BRBeans), review the topic Overview of Business Rule Beans. The following sections provide an overview of externalizing business rules using Business Rule Beans:

1. Developing BRBeans
2. Assembling applications for use with BRBeans
3. Managing rules

To help you get started quickly, a sample BRBeans application is provided. Samples are installed by default during a typical WebSphere installation or you can select to install specific samples during a custom installation.

The BRBeans sample is an online movie store application. The application's EAR file is installed and the application is configured to use the Cloudscape database (Cloudscape is provided with WebSphere Application Server). A number of rules are created that you can view using the Rule Management Application (RMA). To do this, change to the `<install_root>/bin` directory and type one of the following commands:

- On a Microsoft Windows platform:

```
rulemgmt ..\samples\lib\BRBeans\movieSampleProperties
```
- On a Unix platform:

```
rulemgmt.sh ../samples/lib/BRBeans/movieSampleProperties
```

By running the sample, you can see how these rules are used. The source code for the sample also is provided in the `<install_root>/samples/src/BRBeans/Movie` directory. To see the use of trigger points, search the code for places where the `trigger()` method is used.

---

### Advantages of externalizing business rules

Business Rule Beans (BRBeans) provide a framework in which business applications can externalize business rules. You can externalize rules by extending your application analysis and design processes to identify the points of variability (or "trigger points") in application behavior. When the application is implemented, the business logic required at the points of variability is externalized into a business rule. This allows certain aspects of the behavior to be changed without actually changing the application code.

Here are some advantages of externalizing business rules:

#### **Explicit documentation of business practice decisions**

Separating business rule values from the application code makes the code easier for others to view and understand while isolating information that relates to business practice decisions.

#### **Clearer understanding of application behavior**

Externalization makes it possible to inspect the application to see which business rules are being applied, when they are applied, and under what circumstances.

#### **Reuse of rules across business processes**

Separating rules from the business logic of the application makes it easy to reuse a business practice decision in a consistent fashion.

#### **Increased consistency of business practices**

Because externalized rules promote reuse and facilitate clear understanding of business practice decisions, they provide a basis for improving business practice consistency across applications.

#### **Decreased maintenance and testing costs**

Externalized rules have a clearly defined scope and are not tightly coupled to the application code. This makes them easy to modify, quick to test, and decreases costs and improves cycle time.

### **Improved manageability of business practice decisions**

Externalization, change history, and inspectability all promote clear ownership and consequently a better definition of who can change rules and under what circumstances.

### **Increased confidence in predicting the business impact of proposed changes**

Because rules are available for inspection, have well-defined scope, and are not tightly coupled to application business logic, they make it easy to understand the likely impact of changes and to predict whether contemplated modifications or additions will have unwanted ripple effects.

### **Ability to identify and correct conflicting business rules in different parts of the business**

Externalized rules make it easy to check that rules being used in two different parts of an application or even two different applications dealing with different parts of the business, are consistent.

---

## **Overview of Business Rule Beans**

Business Rule Beans are used to create and modify rules that keep pace with complex business practices. This enables your application's core behavior and user interface objects to remain intact and untouched, even as business practices change.

The Business Rule Beans (BRBeans) framework enables you to organize rules in folders. Folders provide a structure similar to the file system on your computer's hard drive. For example:

- Rules can be placed in folders based on any criteria you want.
- A rule folder can contain any number of rules and other rule folders.

In the BRBeans framework, each business rule is represented by an entity bean that persistently stores information related to that rule. Each business rule is assigned an appropriate rule name and stored in an appropriate rule folder (See Rule folders for more information).

When naming rules and folders, adhere to the Java package naming convention. That is, name rules and folders based on the domain name of the organization for which the rules are developed. For example, ACME's `isSeniorCitizen` rule's fully qualified rule name ("full rule name"), might be `com/acme/ageRules/isSeniorCitizen`. In this example, the `com/acme` path is used by all of the rules developed by ACME and the `ageRules` folder is used to separate "age" rules from rules of other kinds. The root folder has no name; therefore, fully qualified path names never start with a forward slash (`/`).

A fully qualified rule name consists of the following:

- The full path of the folder followed by a forward slash (`/`)
- The name of the rule

This fully qualified rule name is used by a trigger point to identify the rule to trigger. Trigger points are small pieces of code that interface with the Business Rule Beans trigger point framework to run business rules during application execution. See [Placing a trigger point](#) for more information.

By default, trigger points can only trigger rules that are currently in effect based on the current date and time when the trigger point is called. A business rule has a start date and an end date (see [Rule attributes](#) for more information) that together define the interval during which the rule is in effect (see [Rule states](#) for more information). This behavior can be overridden by specifying a date on the trigger point. This date is referred to as the "As Of Date". If no start date is specified, the rule is not valid and cannot be found by trigger points. Conversely, if no end date is specified, the rule never expires. Dates and times with a precision of one second can be assigned using the "Rule Management Application".

When there is more than one rule with the same fully qualified name, all of the rules with that name that are currently in effect are triggered and the results are combined using the combining strategy specified on the trigger point. See the `CombiningStrategy` method for more information.



## Externalized business rules

A business rule is a statement that defines or constrains some aspect of a business by asserting control over some behavior of that business.

A business rule officiates over frequently changing business practices and can come from within the company or be mandated from outside, typically by regulatory agencies. Typical uses for business rules include the following:

- Determining the current interest rate for a home loan
- Calculating a discount for a product
- Calculating the tax to apply to a given product
- Determining whether a given person is a senior citizen

The objects used to implement a business rule contain methods and attributes used by the Business Rule Beans (BRBeans) run-time environment, its administrative component, or both. An externalized business rule is implemented as a pair of objects:

- Rule
- RuleImplementor

The **Rule** is an entity enterprise bean that stores all of the persistent data for the business rule. This is the object that the trigger point framework code actually deals with directly. When a trigger point is invoked, the internal framework code performs a query to find the Rule object or objects representing the business rules to be triggered. Once the Rules are found, the framework code determines where the Rule is invoked, either local to the trigger point or remotely on the application server. Then, it invokes the `fire` method on either the Rule enterprise bean itself (for remote triggering) or on a local copy of the enterprise bean (for local triggering) to perform the function of the business rule.

The class name of the business rule's **RuleImplementor** is stored persistently in the Rule. The RuleImplementor is a transient object (not managed by the application server) that the Rule instantiates and then uses to do the actual work. When the `fire()` method is called on the Rule object, the Rule object combines its persistent set of values with the parameters it received on invocation. This creates the parameter list for the RuleImplementor prior to invoking `fire()` on the RuleImplementor with this parameter list. The actual execution of the RuleImplementor algorithm can take place either remotely (within the application server where the BRBeans enterprise beans are installed) or locally (within the Java virtual machine (JVM) where the trigger point was called).

## Types of business rules

Business rules can be divided into the following types:

- Base rules
- Classifier rules

**Base rules** are the most common type of rule and are triggered by the `TriggerPoint.trigger` method.

You can divide Base rules into the following categories:

### Derivation rules

These rules use an algorithm to return a value. These rules return any type of value that makes sense in the business context in which they are used. For example, a derivation rule can calculate a discount or compute the total price of an order.

### Constraint rules

These rules confirm that an operation has met all of its obligations and that a particular constraint or edit has been met. For instance, a constraint rule can check that a value entered by an external user is within legal bounds. Business Rule Beans (BRBeans) provide a special return type, `com.ibm.websphere.brb.ConstraintReturn`, which can be returned by a constraint-type rule. A `ConstraintReturn` object contains a boolean value so that if it is false, it can contain information that can be used to produce an external message explaining what constraint was not met.

## Invariant rules

These rules ensure that multiple changes made by an operation are properly related to one another.

## Script rules

These rules implement "micro-workflow" or electronic performance support. They are small, variable pieces of a business process that provide assistance to end-users to get the most from the application.

On the surface, **classifier rules** are much like base rules. However, classifier rules can be used to determine the ways in which variables are classified by a business. Classifier rules are triggered by the `TriggerPoint.triggerClassifier` method.

A classifier rule is used to compute a classification for a particular business situation. The classification returned is required to be of type `string`. For instance, bank customers may be classified into gold, silver, and bronze categories based on their spending history or the amount of money they have in their account. For more information on this type of rule, refer to Situational trigger point.

## Rule folders

Rule folders are similar to the directories that divide a computer's hard drive in that they split a large number of files into conceptual units. The rule folder adds its path to the fully qualified rule name. Like the directories on a hard drive, a rule folder can contain any number of rules or rule folders.

Although you can name the folders whatever you deem appropriate, it is recommended that you follow the Java package naming convention. That is, base the names on the domain name of the organization where the rules are developed. So, the fully qualified rule name or full rule name of ACME's `isSeniorCitizen` might be `com/acme/ageRules/isSeniorCitizen`. In this example, the `com/acme` path is used by all of the rules developed by ACME and the `ageRules` folder is used to separate "age" rules from rules of other kinds.

**Note:** The root folder has no name, which means that fully qualified path names never start with a '/'.

When using the Rule Management APIs, a rule folder contains instances of `IRules`, which also are referred to as "rules". To begin working with rules, get the root rule folder by using the `getRootFolder` method on `RuleMgmtHelper` class. From the root rule folder you can add, delete, and retrieve folders and rules using methods on this interface.

## Rule attributes

### Rule name

A name for the rule that is appropriate to its business context.

### Rule folder

The folder that contains the rule.

### Start date

This is the date and time at which the rule goes into effect. Prior to this time, it will not be found by trigger points. Together with the end date, the start date defines a period of time during which the rule is effective. A rule that does not have a start date specified is not a valid rule and will not be found by trigger points.

### End date

This is the date and time at which the rule is no longer effective. After this date and time the rule is no longer in effect and will not be found by trigger points. Together with the start date, the end date defines a period of time during which the rule is effective. A rule that does not have an end date specified is valid and will never expire.

**Ready** This indicates whether the rule is ready to be used. Rules that are not marked as ready will not be found by trigger points. This is intended to be an easy way to keep a rule from being used until it is completely defined or to temporarily turn a rule off without having to change the basic rule data such as start and end dates.

**Java Rule Implementor name**

This is the fully package-qualified name of a Java class that implements the BRBeans RuleImplementor interface. The `fire` method of the class performs the function of the rule. Business Rule Beans (BRBeans) provide several predefined rule implementors or you can write your own. See Rule Implementors or Customized rule implementors for more information.

**Initialization parameters**

This is an array of parameters that are passed to the rule implementor to initialize it. Each element in the array can be any object. This also can be referred to as the rule data, which is the external data that may change over time. The initialization parameters defined for a rule are passed directly to the `init` method of the rule implementor when it is instantiated. See "Rule Implementors" for more information on how rule implementors can use initialization parameters.

**Firing parameters**

Normally, firing parameters are simply the parameters passed on the trigger point when a rule is triggered. However, it is allowed to override these parameters by specifying parameters on the rule itself. This is where these overriding parameters are specified.

**Firing location**

This specifies where the rule implementor for this rule is instantiated and run. The following values are allowed:

**Local** This option instantiates the rule implementor and runs it local to the trigger point (in the same JVM as the trigger point call). This is run on the client machine if the trigger point call is done there or on the server if the server part of an application makes a trigger point call. Use this option for the best performance since, once a rule is cached on the client, the entire triggering process can be performed locally without going to the server at all. The main disadvantage of this option is that the class files for the rule implementors need to be available on every client that can trigger rules.

**Remote**

This will instantiate the rule implementor and run it on the application server where the Business Rule Beans enterprise beans are installed. When using this option at least one remote method call always is required to trigger a rule since the trigger takes place on the server. The advantage is that the rule implementor class files only need to be available on the server.

**Anywhere**

This option tries to instantiate and run the rule implementor locally, and, if the class cannot be found, it tries to trigger it remotely.

**Classification**

For classified rules, this is the classification to which the rule applies. This is used when you use a situational trigger. Once a classification is computed for the situational trigger point, rules that apply to that classification are found and triggered. For more information, see Situational trigger point.

**Classifier**

This indicates whether this rule computes a classification. Classification is used for a situational trigger. A classifier rule is used to perform the first step of a situational trigger which computes a classification that is used to find rules to deal with the situation. For more information, see "Situational trigger point".

**Dependent rules**

In many cases, a rule triggers other rules to complete the overall task. These other rules are referred to as dependent rules and can be specified using the dependent rules attribute. For more information, see Dependent rules.

**Business intent**

This is a text description of the intent of this rule from the view point of the business analyst. You can store any text string here.

**Description**

This is a text description of the rule at the programmer's level. You can store any text string here.

### Original requirement

This is a text description of the initial business analyst requirement of this rule. You can use this description to keep track of why this rule was originally created (for example, to keep auditing records). You can store any text string here.

### User-defined data

You can store a user-defined text string here. The format and use of this data is completely determined by the user.

### Primary key

Every rule has a primary key to uniquely identify it in the database where the enterprise beans are stored. Normally, a unique primary key is generated automatically when you create a new rule. However, you can use the rule management APIs to specify your own primary key, if desired. See Rule management APIs for more information.

### Precedence

This is the relative priority of this rule. The default finding strategy uses this value to order the rules found in the database, from lowest to highest, when more than one rule is found for a particular trigger point. Rules are sorted numerically by precedence with the numerically lowest precedence first and the numerically highest precedence last.

## Rule states

Rules can be in any one of the following states at any particular time:

### scheduled

The rule is scheduled to become effective (its start date is in the future) and will not be found by current trigger points.

### in effect

The rule is currently in effect and can be found by trigger points.

### expired

The rule is no longer in effect (the end date is in the past) and will not be found by trigger points.

### invalid

The rule is not correctly defined and will not be found by trigger points.

Typically, only those rules that are "in effect" are found by the Business Rule Beans (BRBeans) run-time environment. This behavior can be overridden by setting an `asOfDate` on the `TriggerPoint` object, which then will execute "as if" the current date is the given date. For more information, see "As Of Date".

When a Rule is first created, it is marked as "ready for use" and is found when firing Rules. If the Rule is not complete and you do not want it to be found by BRBeans, then use either of the following to mark the Rule:

- Use the `setReady(false)` method in the Rule Management APIs
- Use the Rule Management Application to mark the rule as not ready

## Rule results

In general, a rule can return any type of result that makes sense for the business purpose of the rule. The return type on the `fire()` method is `java.lang.Object` so any Java object can be returned, including arrays. You cannot return a Java primitive since the results must be an object. However, you can return the object form of the primitives. For example, you can return a `java.lang.Integer` instead of an `int`. If the rule is fired remotely, the returned value must implement `java.io.Serializable`.

## Dependent rules

When a business rule triggers other business rules as part of its implementation, the rules that are triggered are called **dependent rules** of the first rule. An example is the `RuleAND` rule implementor supplied with Business Rule Beans (BRBeans). It uses two or more dependent rules, each of which is assumed to return a true or false value. When a rule with `RuleAND` as its implementor is triggered, it triggers each of its dependent rules and a logical AND operation is performed on all of the returned results. The result of this AND operation is returned as the result of the top-level rule.

Dependent rules are specified in the attributes of the top-level rule where the fully qualified name of each dependent rule is listed. When the top-level rule is triggered, an array of dependent rule names is passed to the rule implementor's `init()` method. They are stored here until they are triggered by the `fire()` method.

**Note:** The BRBeans framework does not ensure that the dependent rules specified in the enterprise beans are actually triggered. Triggering the dependent rules and interpreting their results is entirely up to the rule implementor of the top-level rule.

Dependent rules can be nested within other dependent rules. In other words, a dependent rule of some particular rule can have its own dependent rules which, in turn, can have their own dependent rule and so on. The BRBeans framework does not place any restriction on the number of levels that dependent rules can be nested. The only practical restriction is the complexity of the rule set that is built up when dependent rules are nested many levels deep.

## BRBeans run-time environment

The Business Rule Beans (BRBeans) run-time environment is used to find and trigger rules.

The BRBeans run-time environment is made up of two parts:

- Code that runs on the **client** ("client" here meaning wherever the trigger point is located). This consists of code that does the following:
  - Finds the specified rules
  - Decides where the rules should be triggered
  - Calls the `fire` method on all rules
  - Combines the results from the rules
- Code that runs on the **server**. This consists of enterprise beans used to represent rules and rule folders. These enterprise beans do the following:
  - Provides for business rule persistence
  - Provides query functions that the client part of the run time can use to find rules to be triggered

## BRBeans run-time behavior

Business Rule Beans (BRBeans) run-time behavior can best be described by giving a simple example of a trigger point selecting, executing, and then responding to the results of a business rule.

The first step in triggering a rule is for the trigger point framework to perform a query on the rule server to determine which rules to trigger. The main item used for the query is the fully qualified rule name. Other items used in the query include the start and end date, whether this is a classifier, the classification of the rule, and whether the rule is marked "ready". This query returns zero or more rules. If there is at least one rule, the trigger point assembles the data that is sent as parameters to each rule. The trigger point then loops through the list of rules invoking the `fire()` method on each and passing the parameters. The results are combined depending on the combining strategy used.

When the trigger point framework invokes `fire` on a rule, it instantiates the `RuleImplementor` and uses it to do the actual work (to execute the rule algorithm or test). Once it has arrived at a result, the `RuleImplementor` returns that result. For constraint rules (ones that arrive at a boolean true or false answer) the returned value is, by convention, a `ConstraintReturn`. A `ConstraintReturn` is a data structure that indicates whether the constraint was satisfied. If not, the `ConstraintReturn` indicates what went wrong. For derivation rules (ones that calculate a single, generally non-boolean value), the return value can be any type. In the simplest case, the return value from each `RuleImplementor` is returned back to the trigger point where it is analyzed to determine what action to take.

The following is an overview of what happens when the `maxTruckDriverHours` rule is triggered:

A rule exists named `maxTruckDriverHours`. The purpose of this rule is to check that the number of hours entered by a user for a particular truck driver does not exceed the maximum allowed value. This rule contains an initialization parameter list consisting of a single value of 8. This rule is bound to a

RuleImplementor class called MaxRuleImpl. MaxRuleImpl tests the parameter it is passed against the initialization list value and returns a ConstraintReturn. The ConstraintReturn is set to true if the passed parameter is less than or equal to the initialization value. Otherwise, a ConstraintReturn is set to false and some information is added that describes which values were compared and why the test failed.

When this rule is triggered, the following details the trigger point process:

1. During the execution of the application, the application reaches a point where it needs to verify that the number of truck driver hours that was entered is valid. The application code invokes a simple trigger point passing the name of the rule to be triggered and a parameter list containing the entered hours for the driver.
2. The trigger point framework performs a query on the rule server to find the rule with the specified name. It receives back a sequence of rule objects. In this case, this sequence contains one rule, maxTruckDriverHours.
3. The framework determines whether this rule is to be triggered on a local or remote machine. If local, the framework gets a local copy of the rule object and calls the fire method on the copy. If remote, the framework calls the fire method on the enterprise bean reference. The parameter list containing the weight is passed to the fire method.
4. The maxTruckDriverHours rule (either the copy or the enterprise bean itself) creates an instance of the rule implementor class, maxRuleImpl, if it does not already have one. When a new rule implementor instance is created, the rule calls its init method passing any initialization parameters defined for the rule. In this case, the initialization parameter list contains the single value 8. If the rule already has a rule implementor instance, it uses that one and does not call the init method.
5. The maxTruckDriverHours rule calls the fire method on the rule implementor instance. The firing parameters passed to the trigger point are passed to the rule implementor and are possibly modified by any firing parameters defined in the rule itself. In this case, the firing parameters are passed directly from the trigger point.
6. The maxRuleImpl returns a ConstraintReturn object to the rule that indicates the result of its comparison. This ConstraintReturn is returned to the trigger point framework and ultimately to the application.
7. The application checks the value in the ConstraintReturn and takes the appropriate action.

## BRBeans run-time exception handling

Business Rule Beans (BRBeans) defines one general exception class for exceptions that might be exposed to the user. All other BRBeans exceptions inherit from this class. The name of this class is `com.ibm.websphere.brb.BusinessRuleBeansException`. A `BusinessRuleBeansException` is generally thrown when an unexpected error occurs within BRBeans. A `BusinessRuleBeansException` might have information in it about the original exception that caused the error. Doing a `printStackTrace` on the `BusinessRuleBeansException` prints out this information and the stack trace for the `BusinessRuleBeansException` itself. Also, there are methods to access the original exception programmatically, if desired.

BRBeans also defines a **ConstraintViolationException**, which extends `BusinessRuleBeansException`. A `ConstraintViolationException` is thrown if the `ThrowViolationCombiningStrategy` is specified on the `TriggerPoint` and the rule returns a false value (either a `ConstraintReturn` or a boolean).

Finally, BRBeans defines two exceptions, **NoRuleFoundException** and **MultipleRulesFoundException**, that are thrown by some of the predefined filtering strategies if an unexpected number of rules is found on a trigger point call. These two exceptions both extend `UnexpectedRulesFoundException` which, in turn, extends `BusinessRuleBeansException`.

## Rule implementors

A rule implementor, in terms of Business Rule Beans, is an algorithm written in Java that implements the `BRBeans RuleImplementor` interface.



A Business Rule Beans (BRBeans) Rule is a persistent object that exists on the BRBeans Rule server. One of the rule's persistent attributes, in addition to `startDate`, `endDate`, `initParams`, and so on, is `javaRuleImplementorName`, which is the name of its rule implementor.

BRBeans supplies a number of predefined rule implementor classes that can be used in user-defined BRBeans rules (see the BRBeans Javadoc) to implement the BRBeans `RuleImplementor` Interface. The Java source code for these rule implementors is supplied as BRBeans sample code in the `com.ibm.websphere.brb.implementor` package. This sample code, packaged in a JAR file, appears in the CLASSPATH of the BRBeans Rule Server (for "remote" firing) or is co-located in the CLASSPATH of the application or applications using it (for "local" firing). Typically, the `RuleImplementor` is in the application EAR file.

Using standard Java development tools, you can externalize BRBeans by attaching them to either enterprise beans or ordinary Java objects. Programming a new rule implementation in Java is typically a simple process. If you write your own rule implementor, you must create a new Java class that implements the `com.ibm.websphere.brb.RuleImplementor` interface. This class must implement the following methods:

**Default constructor**

The class must have a default, no-argument constructor so that it can be instantiated when a rule using it is triggered.

**init** The `init` method comes from the `RuleImplementor` interface and is called when the rule implementor is first. Its purpose is to perform an initialization needed by the rule implementor instance before it is actually fired. The following parameters are passed to the `init` method:

**The initialization parameters defined for the rule being triggered**

These can be any parameters needed to properly initialize the rule implementor instance. Often the initialization parameters consist of constants required by the algorithm. For example, when using a rule implementor that checks whether a number is greater than a threshold value, the threshold value normally is passed as an initialization parameter. This parameter is null if there are no initialization parameters for the rule.

**An array of names of dependent rules for the rule being triggered**

Normally, the rule implementor stores these names to be used when the `fire` method is called. These dependent rules are intended to be triggered as part of the algorithm performed by the rule implementor. See *Dependent Rules* for more information. This parameter is null if there are no dependent rules defined for the rule.

**The user-defined data for the rule being triggered**

This data is completely defined by the user of the Business Rule Beans (BRBeans). BRBeans does not interpret this data in any way. This parameter is null if there is no user-defined data defined for the rule.

**A reference to the actual rule being triggered**

This can be used to extract attribute values from the rule, if needed.

**fire** The `fire` method comes from the `RuleImplementor` interface. This method is called to perform the algorithm of the rule implementor. Any desired algorithm can be performed here. Normally, a value is returned by the `fire` method that is ultimately returned as the result of triggering the rule. The following parameters are passed to the `fire` method:

**The TriggerPoint object that is being used to trigger the rule**

This parameter is needed if the rule has dependent rules that the `fire` method needs to trigger.

**The target object for this particular trigger call**

This parameter can be any object that is thought of as the target of the rule. However, the parameter can be null.

**A reference to the actual rule being triggered**

This parameter can be used to extract attribute values from the rule, if needed.

**The firing parameters for this particular trigger call**

Normally, these parameters are the firing parameters passed by the code that invoke the trigger point. However, these can be overridden by specifying firing parameters on the rule itself. Wherever they ultimately come from, these are the parameters that the rule implementor needs at run time to perform its function. Normally, these are run-time



variables that are to be processed by the rule implementor. For example, when using a rule implementor that checks whether a number is greater than a threshold value, the number to be checked normally is passed as a firing parameter. This parameter is null if no firing parameters are passed by the caller and none are defined on the rule itself.

### **getDescription**

`getDescription` comes from the `RuleImplementor` interface. The purpose of this method is to return a text string that describes the function of the rule implementor. This information might be displayed on a user interface to help select what implementor to use. This method, however, is currently not used by the BRBeans framework. Users can incorporate this information if they create their own rule implementor. For additional information, see the BRBeans Javadoc for the `RuleImplementor` interface.

## **Trigger point framework**

A trigger point is the location in a method of an object where externalized business rules are invoked.

The proper placement of trigger points can add substantially to the flexibility and speed with which a business application adapts to new business practices.

Wherever a trigger point is placed in user-written code, the Business Rule Beans (BRBeans) trigger point framework needs to do the following:

1. Assemble the parameter list to send to the rules
2. Find the potential rules that apply
3. Filter out any rules which do not apply (optional)
4. Fire the rules in the filtered rule set
5. Combine the results of the rule firings in some meaningful way

The application code that contains the trigger point needs to perform the following functions:

1. Establish a value for the target object. Usually the target object is the object in which the trigger point is encountered. The target object is one of the parameters passed to the `fire` method of the `RuleImplementor`. However, this is an optional parameter. If the rule implementor does require a target object, null can be passed instead.
2. Build the array of objects containing the run-time parameters needed to satisfy the trigger point's business purpose. This array is normally passed as one of the parameters of the `fire` method of the `RuleImplementor`. If firing parameters are specified on the rule itself, then those firing parameters are passed instead of the ones passed by the caller.
3. Invoke the `trigger()`, `triggerClassifier()`, or `triggerSituational()` method of the `TriggerPoint` class.
4. Catch and handle any exceptions that might occur as a result of firing the rules. Otherwise, take action based upon the rule firing results.

The two simple trigger methods, `trigger` and `triggerClassifier`, perform their function in four steps:

1. Find the rules
2. Filter out those rules which are not desired
3. Fire the remaining rules
4. Combine the results and return to the caller

The complex trigger method, `triggerSituational` does this sequence of steps twice. In the first phase, the method performs the four steps once to find a rule that returns a classification. This classification is fed into the second phase. The second phase triggers rules that have the name specified in the `triggerSituational` method and have a classification equal to the value returned by the first phase.

How each of these steps is performed can be modified through various methods on the `TriggerPoint` object. The implementation of each step is defined by a **strategy** object. For more information on strategies, see *Administering strategy objects to control triggers*.

## Trigger points

Examples of how to code a trigger point call are provided in the following topics:

- Simple trigger point
- Classifier trigger point
- Situational trigger point

### Simple trigger point

A simple trigger point is used to trigger a rule or rules specified by name. This type of trigger point is used by invoking the trigger method on an instance of the `TriggerPoint` class. All rules with the specified name are triggered and the results are combined using the `CombiningStrategy` specified on the `TriggerPoint` object. This type of trigger point only finds rules that are not marked as classifiers.

The following shows an example of using a simple trigger point to trigger a rule named `isSeniorCitizen` (in the `com/acme/ageRules` folder), which determines whether a person is classified as a senior citizen based on the passed in age:

```
...
// create an instance of TriggerPoint for triggering the rule and specify that the
// ReturnFirstCombiningStrategy is to be used to return only the first result if
// multiple rules are found.
TriggerPoint tp = new TriggerPoint();
tp.setCombiningStrategy(CombiningStrategy.RETURN_FIRST, TriggerPoint.ALL_RULES);

// define parameter list that's passed to the rule
Object [] plist = new Object[1];

// define age of person to be tested
Integer age = new Integer(64);

// define name of rule to be fired
String ruleName = "com/acme/ageRules/isSeniorCitizen";

// define result of rule firing
Object result = null;

// initialize parameter list
plist[0] = age;

try {

 // fire "com/acme/ageRules/isSeniorCitizen" rule passing parameter list containing age.
 // Note: in this case the target object is not used and could be null.
 result = tp.trigger(this, plist, ruleName);

 // put result into usable format. A single result is returned since we specified to use
 // the ReturnFirstCombiningStrategy. By default an array of results would be returned.
 boolean seniorCitizen = ((Boolean)result).booleanValue();

 // make use of result
 if(seniorCitizen) {
 ...
 }
}
catch(BusinessRuleBeansException e) {

 // handle exception
 ...
}
}
```

## Classifier trigger point

A classifier trigger point is identical to a simple trigger point except that it only finds rules marked as classifiers. Classifiers are rules that determine what sort of business situation is present. These rules then return a classification string that indicates the result.

Usually these rules are used as part of a situational trigger point, but they also can be triggered on their own. This type of trigger point is used by invoking the `triggerClassifier` method on an instance of the `TriggerPoint` class.

The following shows an example of using a classifier trigger point to trigger a rule named `determineCustomerLevel` (in folder `com/acme/customerClassifiers`). This rule classifies customers into levels (gold, silver, and bronze) based on their spending history.

```
...
// create an instance of TriggerPoint for triggering the rule and specify that the
// ReturnFirstCombiningStrategy is to be used to return only the first result if
// multiple rules are found.
TriggerPoint tp = new TriggerPoint();
tp.setCombiningStrategy(CombiningStrategy.RETURN_FIRST, TriggerPoint.ALL_RULES);

// define parameter list that's passed to the rule
Object [] plist = new Object[1];

// information about the customer to be checked is stored in this object
Customer cust = ...;

// define name of rule to be fired
String ruleName = "com/acme/customerClassifiers/determineCustomerLevel";

// define result of rule firing
Object result = null;

// initialize parameter list
plist[0] = cust;

try {

 // fire "com/acme/customerClassifiers/determineCustomerLevel" rule passing parameter
 // list containing the customer to be checked.
 // Note: in this case the target object is not used and could be null.
 result = tp.triggerClassifier(this, plist, ruleName);

 // put result into usable format. A single result is returned since we specified to use
 // the ReturnFirstCombiningStrategy. By default an array of results would be returned.
 String customerLevel = (String) result;

 // make use of result
 if(customerLevel.equals("Gold")) {
 ...
 } else if (customerLevel.equals("Silver")) {
 ...
 } else if (customerLevel.equals("Bronze")) {
 ...
 } else {
 ...
 }
}
catch(BusinessRuleBeansException e) {

 // handle exception
 ...
}
}
```

## Situational trigger point

A situational trigger point is used when the rule or rules to be triggered depend on the business situation.

This example evaluates a customer's past purchasing history to place them into one of three levels: Gold, Silver, or Bronze. Their classification determines how much of a discount they receive.

To use a situational trigger point to handle this case, it is first necessary to define four rules:

- one **classifier** rule to determine under which of the three levels to classify the customer
- three **classified** rules to determine the actual discount to offer

All of the classified rules have the same name and are marked as applying to one of the three customer levels by specifying the level in its classification attribute. For example, the rule to determine the discount for a Gold level customer will contain the string "Gold" in its classification attribute.

The situational trigger point takes two rule names as input: the name of the classifier rule and the name of the classified rule. The situational trigger point then proceeds in two phases:

1. Find the specified classifier rule and trigger it to generate a classification string.
2. Find the rules that have the name specified for the classified rule and have a classification attribute equal to the classification string returned by the first phase.

These rules then are triggered to produce the final result, in this case the discount to offer.

The following shows an example of a situational trigger point used to handle the case described previously:

```
...
// create an instance of TriggerPoint for triggering the rule and specify that the
// ReturnFirstCombiningStrategy is to be used to return only the first result if
// multiple rules are found.
TriggerPoint tp = new TriggerPoint();
tp.setCombiningStrategy(CombiningStrategy.RETURN_FIRST, TriggerPoint.ALL_RULES);

// define parameter list that's passed to the classifier rule
Object [] classifierPlist = new Object[1];

// define parameter list that's passed to the classified rule
Object [] classifiedPlist = new Object[1];

// information about the customer to be checked is stored in this object
Customer cust = ...;

// define name of classifier rule to be fired
String classifierRuleName = "com/acme/customerClassifiers/determineCustomerLevel";

// define name of classified rule to be fired
String classifiedRuleName = "com/acme/discountRules/determinedDiscount";

// define result of rule firing
Object result = null;

// initialize parameter lists
classifierPlist[0] = cust;
classifiedPlist[0] = cust;

try {
 // fire the rules to get the discount to offer
 // Note: in this case the target object is not used and could be null.
 result = tp.triggerSituational(this, classifiedPlist, classifierPlist, classifiedRuleName, classifierRuleName);

 // put result into usable format. A single result is returned since we specified to use
 // the ReturnFirstCombiningStrategy. By default an array of results would be returned.
 Float discountToOffer = (Float) result;
}
```

```

// make use of result
...
}
catch(BusinessRuleBeansException e) {

// handle exception
...
}

```

## As Of Date

An "As Of Date" can cause rules to be triggered as if the given date is the current date. This is especially useful when you want to test a rule, see what effect a future change in rules or regulations may have on the overall framework, or see what past or future rates, discounts, or both might be.

Normally, a rule only can be triggered if it is "in effect" (see "Rule States") as of the current date and time.

To set an "As Of Date", call the `setAsOfDate()` method on the `TriggerPoint` object and pass the date that you want to be used. To use the current date again, call `unsetAsOfDate` or `setAsOfDate` and pass null for the date.

## Predefined strategy objects

The following is a list of predefined strategy objects that are provided in Business Rule Beans:

### FindingStrategy

Accesses the data store and returns those rules that meet the search criteria specified

### FilteringStrategy

Takes the list of rules that were found by the `FindingStrategy` and filters out the rules that should not be fired

### FiringStrategy

Takes the rules that were found by the `FindingStrategy`, (possibly modified by the `FilteringStrategy`), fires them each in order, and returns an array containing the results of each rule

### CombiningStrategy

Takes the results of the rules that are fired by the `FiringStrategy` and combines them to form a reasonable result to the `TriggerPoint` caller.

## FindingStrategy method

The job of the `FindingStrategy` is to access the data store and return those rules that meet the search criteria specified. There are two `FindingStrategy` classes provided by Business Rule Beans (BRBeans):

- `DefaultClassifierFindingStrategy`
- `DefaultNonClassifierFindingStrategy`

Both of these strategies perform a case-sensitive search for Rules that are marked "ready" and match the given search criteria. Results are ordered by precedence from highest to lowest (the first rule in the array has the numerically smallest precedence, the next rule has the next smallest precedence, and so on). If no rules are found, then an empty array is returned. The former strategy returns classifier rules (`classifier=true`) only and the latter returns non-classifier rules (`classifier=false`) only.

These default strategies are used automatically by the `TriggerPoint`. There is no need to call `setFindingStrategy` to use these strategies. Instances of these two default finding strategies are stored in static constants defined on the `FindingStrategy` interface.

## FilteringStrategy method

The job of the `FilteringStrategy` is to take the list of rules that were found by the `FindingStrategy` and filter out the rules that should not be fired. There are three sets of filtering strategies used in

`TriggerPoint`:

- strategy for **zero** rules found
- strategy for **one** rule found

- strategy for **multiple** rules found

A different strategy can be used for each of these scenarios, along with different strategies for classifier and non-classifier rules. The zero rules strategy is invoked if no rules are found by the finding strategy, the one rule strategy is invoked if exactly one rule is found and the multiple rules strategy is invoked if more than one rule is found.

Business Rule Beans (BRBeans) provides the following filtering strategies that can be used:

**Accept Any**

BRBeans utilizes all of the rules found (this is the default).

**Accept One**

BRBeans expects one rule only.

**Accept First**

BRBeans utilizes the first rule found.

**Accept Last**

BRBeans utilizes the last rule found.

Instances of these filtering strategies are stored in static constants defined in the `FilteringStrategy` interface. You can use these for setting the strategies on a `TriggerPoint`.

As an example, here is one common way to use filtering strategies. You want to ensure that exactly one rule is found on a `TriggerPoint` call. Thus, set all three strategies (zero rules, one rule, and multiple rules) for this `TriggerPoint` to `FilteringStrategy.ACCEPT_ONE`. This strategy throws an exception if the number of rules is not exactly one. The following sequence of method calls accomplishes this for `TriggerPoint tp`:

```
tp.setNoRulesFilteringStrategy(FilteringStrategy.ACCEPT_ONE, TriggerPoint.ALL_RULES);
tp.setOneRuleFilteringStrategy(FilteringStrategy.ACCEPT_ONE, TriggerPoint.ALL_RULES);
tp.setMultipleRulesFilteringStrategy(FilteringStrategy.ACCEPT_ONE, TriggerPoint.ALL_RULES);
```

## FiringStrategy method

The `FiringStrategy` takes the rules that were found by the `FindingStrategy`, (possibly modified by the `FilteringStrategy`), fires them each in order, and returns an array containing the results of each rule.

A single default `FiringStrategy` is provided by Business Rule Beans (BRBeans) as all of the rules are fired in the same way. This implementation takes each rule in order and performs the following steps:

1. Determines what firing parameters to pass to the rule. If there are no firing parameters specified for this rule, the implementation uses the firing parameters passed on the `TriggerPoint` call. Otherwise, it uses the firing parameters specified in the rule in place of the parameters passed on the `TriggerPoint` call.
2. Calls the `fire` method on the rule and passes the firing parameters from the first step.

Unexpected exceptions result in a `BusinessRuleBeansException` being thrown that contains the original exception.

## CombiningStrategy method

The job of the `CombiningStrategy` is to take the results of the rules that are fired by the `FiringStrategy` and combine them to form a reasonable result to the `TriggerPoint` caller. Business Rule Beans (BRBeans) provides several combining strategies to be used in applications:

**Return All**

Returns the results from all of the rules fired in an array (this is the default)

**Return First**

Returns only the result from the first rule fired

**Return Last**

Returns only the result from the last rule fired

**Return AND**

Returns the logical AND of the results from all the rules fired. This strategy requires that all of the

results returned by the fired rules are either `ConstraintReturn` objects or `java.lang.Boolean` objects. An exception is thrown if this is not the case.

#### **Return OR**

Returns the logical OR of the results from all of the rules fired. This strategy requires that all of the results returned by the fired rules are either `ConstraintReturn` objects or `java.lang.Boolean` objects. An exception is thrown if this is not the case.

#### **Throw Violation**

Throws a `ConstraintViolationException` containing all of the failed `ConstraintReturn` objects if any `ConstraintReturns` contain `false`. Otherwise, it returns a `true` `ConstraintReturn`.

Instances of these combining strategies are stored in static constants defined in the `CombiningStrategy` interface. You can use these for setting the strategies on a `TriggerPoint`. For example, the following method call sets the `CombiningStrategy` on `TriggerPoint` `tp` to be the `Return_First` strategy:

```
tp.setCombiningStrategy(CombiningStrategy.RETURN_FIRST, TriggerPoint.ALL_RULES);
```

## **Customized strategy objects**

The process of triggering a rule or set of rules is controlled by a set of strategy objects. The following four strategies are used each time a rule is triggered:

#### **FindingStrategy**

The `FindingStrategy` accesses the persistent data store to find the set of rules matching the search criteria passed to the trigger call. The search criteria are based on the rule ID information passed on the trigger call. The set of rules found is passed to the `FilteringStrategy`.

#### **FilteringStrategy**

The `FilteringStrategy` can change the set of rules that were found by the `FindingStrategy`. The set of rules returned is the set that are fired by the `FiringStrategy`.

#### **FiringStrategy**

The `FiringStrategy` fires the rules found by the `FindingStrategy`, which may be modified by the `FilteringStrategy`. It gathers the results of the individual rules and passes them to the `CombiningStrategy`.

#### **CombiningStrategy**

The `CombiningStrategy` takes the results from firing the rules and combines them to produce the final result of the trigger.

Each `TriggerPoint` object has its own set of strategies that can be changed independent of any other `TriggerPoint` object. There is a set of default strategies that are used by the `TriggerPoint` if none are explicitly set.

For each of the four strategies, you can set different strategies for **classifier** rules and for **non-classifier** rules. The strategies set for classifier rules are used when the Business Rule Beans (BRBeans) framework is triggering a classifier rule. The strategies for non-classifier rules are used in all other cases.

It is also possible to set three different sets of filtering strategies:

- one to be used if no rules are found
- one to be used if exactly one rule is found
- one to be used if more than one rule is found

This capability can be used to set up filtering strategies that throw exceptions if the expected number of rules is not found.

Strategy classes must implement one of the strategy interfaces provided by BRBeans in the `com.ibm.websphere.brb` package:

- `FindingStrategy`
- `FilteringStrategy`
- `FiringStrategy`
- `CombiningStrategy`



Users can write their own strategy implementations to perform special functions not performed by the predefined implementations. Write these strategy implementations with care since part of the functionality of the BRBeans framework is replaced when you write a custom strategy. One simple example of writing a custom strategy is creating a new firing strategy that logs every rule that is fired.

The basic requirement for a strategy implementation is that it implements the appropriate strategy interface.

For the **filtering** and **combining** strategies, create a class that implements either `FilteringStrategy` or `CombiningStrategy` and either the `filterRules()` method (for `FilteringStrategy`) or the `combineResults()` method (for `CombiningStrategy`) to perform the required functions. At run time, create an instance of the new class and pass it to the `TriggerPoint` object using the appropriate set method so that the new strategy is used when rules are triggered using that `TriggerPoint`.

The **finding** and **firing** strategies are more complicated to customize since they provide more function than the simple filtering and combining strategies. Default finding and firing strategy implementations are provided that define a general outline of the steps necessary to perform the function. It is suggested that you subclass these when you customize your own strategies and then override the desired methods on the default implementation to provide the new behavior.

The BRBeans Javadoc provides more information about the `FindingStrategy`, `FilteringStrategy`, `FiringStrategy`, and the `CombiningStrategy` interfaces.

## Customized rule implementors

To write your own rule implementor, create a new Java class that implements the `com.ibm.websphere.brb.RuleImplementor` interface. This class must implement the following methods:

### Default constructor

The class must have a default, no-argument constructor so that it can be instantiated when a rule using it is triggered.

**init** The `init` method comes from the `RuleImplementor` interface and is called when the rule implementor is first created. Its purpose is to perform an initialization needed by the rule implementor instance before it is actually fired. The following parameters are passed to the `init` method:

#### The initialization parameters defined for the rule being triggered

These can be any parameters needed to properly initialize the rule implementor instance. Often the initialization parameters consist of constants required by the algorithm. For example, when using a rule implementor that checks whether a number is greater than a threshold value, the threshold value normally is passed as an initialization parameter. This parameter is null if there are no initialization parameters for the rule.

#### An array of names of dependent rules for the rule being triggered

Normally, the rule implementor stores these names to be used when the `fire` method is called. These dependent rules are intended to be triggered as part of the algorithm performed by the rule implementor. See `Dependent Rules` for more information. This parameter is null if there are no dependent rules defined for the rule.

#### The user-defined data for the rule being triggered

This data is completely defined by the user of the Business Rule Beans (BRBeans). BRBeans does not interpret this data in any way. This parameter is null if there is no user-defined data defined for the rule.

#### A reference to the actual rule being triggered

This can be used to extract attribute values from the rule, if needed.

**fire** The `fire` method comes from the `RuleImplementor` interface. This method is called to perform the algorithm of the rule implementor. Any desired algorithm can be performed here. Normally, a value is returned by the `fire` method that is ultimately returned as the result of triggering the rule. The following parameters are passed to the `fire` method:

**The TriggerPoint object that is being used to trigger the rule**

This parameter is needed if the rule has dependent rules that the fire method needs to trigger.

**The target object for this particular trigger call**

This parameter can be any object that is thought of as the target of the rule. However, the parameter can be null.

**A reference to the actual rule being triggered**

This parameter can be used to extract attribute values from the rule, if needed.

**The firing parameters for this particular trigger call**

Normally, these parameters are the firing parameters passed by the code invoking the trigger point. However, these can be overridden by specifying firing parameters on the rule itself. Wherever they ultimately come from, these are the parameters that the rule implementor needs at run time to perform its function. Normally, these will be run-time variables that are to be processed by the rule implementor. For example, when using a rule implementor that checks whether a number is greater than a threshold value, the number to be checked normally is passed as a firing parameter. This parameter is null if no firing parameters are passed by the caller and none are defined on the rule itself.

**getDescription**

getDescription comes from the RuleImplementor interface. The purpose of this method is to return a text string that describes the function of the rule implementor. This information might be displayed on a user interface to help a user select what implementor to use. This method, however, is currently not used by the BRBeans framework. For additional information, see the RuleImplementor interface in the BRBeans Javadoc.

## Rule management command

The Rule management command assists the user in performing high-level administration of rules and rule folders.

This includes the capability to create, modify, delete, import, or export rules or rule folders. This command can be used initially by the programmer to define rules interactively and then used by the domain analyst for rule management tasks. You can use the following files:

- On Microsoft Windows platforms, rulemgmt.bat
- On Unix platforms, rulemgmt.sh

rulemgmt properties-file [host-address port-number]

**Parameters****<properties-file>**

The fully qualified name of a file containing the JNDI names of the rule EJBs for the rule set that is to be accessed.

The following must be specified in the file:

```
RuleJndi=<JNDI name of the Rule EJB>
RuleFolderJndi=<JNDI name of the RuleFolder EJB>
RuleHelperJndi=<JNDI name of the RuleHelper EJB>
```

See <WAS\_HOME>/bin/brbeansDefaultProperties for an example.

**Arguments****host-address**

This is the host name of the name server. The default is the local host.

**port-number**

This is the port number of the name server. The default is 2809.

## Rule importer command

The rule importer command imports rules into a database from one or more XML documents.

The rule importer command can be invoked using the Rule Management Application (RMA). The user interface in RMA provides some assistance in specifying the parameters required by the importer. Alternatively, the rule importer can be invoked from the command line using the following scripts:

- For Microsoft Windows platforms, `ruleimporter.bat`
- For UNIX platforms, `ruleimporter.sh`

```
ruleimporter <properties-file> <import-files> [options]
```

### Parameters

#### <properties-file>

The fully qualified name of a file containing the JNDI names of the Business Rule Beans (BRBeans) enterprise beans for the rule set that is to be accessed. Refer to BRBeans Properties file for a definition of the contents of this file. This parameter is required.

#### <import-files>

One or more fully qualified names of the files containing XML rule definitions to be imported. These files must contain XML in the format defined in `<WAS_HOME>\bin\brb.dtd` file. This parameter is required.

### Options

#### -[v]erbose

Shows verbose output while importing. This shows the rule definition of every rule that is imported.

-[t]est Parses the input files only and does not create rules on the application server. This will ensure that there are no errors in the syntax of the rule definitions provided in the XML document. Combined with the **-verbose** option, it also can be used to see exactly what rules will be imported.

#### -[u]pdate

Update the existing rule with values from the input file when a rule in an input file has the same primary key as an existing rule. If this option is not specified, then any rule with the same primary key as an existing rule causes an error and that rule is not imported.

#### -[c]ommitteach

Performs a commit after each rule is created rather than creating all of the rules in a single transaction. If this option is not specified, then all rules are created in a single transaction. This means that if any rule causes an error, the entire transaction is rolled back and none of the rules are imported. If **-commiteach** is specified and a rule causes an error, only that rule is not imported. Other rules are still imported.

#### -[h]ost <host-name>

Specifies the name of the host for the name server. The default is the local host.

#### -[p]ort <port-number>

Specifies the port number for the name server. The default is 2809.

## Rule exporter command

The rule exporter command exports rules from a database into an XML document.

The rules that are exported are determined by an XML document, which is provided to the command. The rule exporter function can be invoked using the Rule Management Application (RMA). The user interface in RMA provides some assistance in specifying the parameters required by the exporter. Alternatively, the rule exporter can be invoked from the command line using the following scripts:

- For Microsoft Windows platforms, `ruleexporter.bat`
- For UNIX platforms, `ruleexporter.sh`

```
ruleexporter <properties-file> <export-list-files> [options]
```

### Parameters

#### <properties-file>

The fully qualified name of a file containing the JNDI names of the BRBeans enterprise beans for the rule set that is to be accessed. Refer to BRBeans Properties file for a definition of the contents of this file. This parameter is required.

### <export-list-files>

One or more fully qualified names of files containing a list of rules, folders, or both to be exported. These files must contain XML in the format defined in the <WAS\_HOME>\AppServer\bin\brb-export-list.dtd file. This parameter is required.

### Options

#### -[v]erbose

Shows verbose output while exporting.

#### -[o]utput <file-name>

Specifies the name of the output file where the XML rule definitions are stored. This is a required parameter.

#### -[h]ost <host-name>

Specifies the name of the host for the name server. The default is the local host.

#### -[p]ort <port-number>

Specifies the port number for the name server. The default is 2809.

## BRBeans properties file

Applications that use the Business Rule Beans (BRBeans) enterprise beans (this includes applications that trigger rules or use the rule management APIs) must specify the JNDI names for these enterprise beans so that the application can find them at run time. If the application is running in a J2EE client container, in a servlet, or on the application server itself (for example, as part of another enterprise bean), then these names probably have been specified by the person who configured the application. If the application is not running in a container, the names must be specified some other way. The BRBeans properties file provides a way to do this.

At run time, the BRBeans code looks for a special Java property that identifies the name of the properties file. This Java property can be specified on the command line as `-DbrbPropertiesFile=<file_name>`. The file specified is expected to contain the JNDI names used to find the BRBeans enterprise beans. The BRBeans framework uses these names when it needs to locate the enterprise beans.

When an application attempts to reference BRBeans enterprise beans, the code first looks for the `brbPropertiesFile` Java property. If this property is specified, the names listed in that file are used to find the enterprise beans and to override any EJB references that were specified in the container (if the application is running in a container). If the property is not specified, then BRBeans attempts to use the EJB references specified in the container.

The host name and port number used to access the name server also can be set in this file. If these are not specified, the BRBeans framework uses the name server used by the container in which the application is running. If the application is not running in a container, then localhost is used for the host name and 2809 is used for the port number.

The properties file must be in the following format (entries can be specified in any order):

```
host=<host-name-for-name-server>
port=<port-number-for-name-server>
RuleJndi=<JNDI-name-for-Rule-EJB>
RuleFolderJndi=<JNDI-name-for-RuleFolder-EJB>
RuleHelperJndi=<JNDI-name-for-RuleHelper-EJB>
```

### Location

A default properties file is shipped as `<WAS_HOME>\AppServer\bin\brbeansDefaultProperties`.

There are a set of JAR files that conform to the BRBeans<database-type>.jar naming convention (depending on the type of database that you want to use). If the JAR files are used without changing the JNDI names, then you also can use the default properties file.

## Usage note

The file name still must be specified even if you want to use the default file. There is no file that is used automatically if the `brbPropertiesFile` property is not set.

The tools shipped with BRBeans (the Rule Management Application, the rule importer, and the rule exporter) all run outside of any container. Hence, the JNDI names need to be specified when these tools are run. The scripts for these tools all require that you pass a properties file name as a command line parameter. This name then is specified as the value for the `brbPropertiesFile` property when the tool is run.

## Database considerations for BRBeans

The following relational databases are supported by Business Rule Beans (BRBeans):

- IBM DB2
- IBM Cloudscape
- Microsoft SQL Server
- Oracle
- Sybase
- Informix

This documentation does not provide you with specific instructions on how to use any of these databases. For help with specific commands, consult the documentation that accompanied your database software. The following are general considerations for relational databases that are supported by BRBeans:

### Large character data

There are several attributes in the BRBeans Rule enterprise beans that might contain large amounts of data. This includes fields such as: `businessIntent`, `dependentRules`, `description`, `firingParameters`, `initParameters`, `originalReq`, and `userDefinedData`. The value for these attributes is stored in a character type column within a database table. Whenever possible, the values are stored in large character fields like `LONG VARCHAR` (for DB2) and `TEXT` (for Sybase).

There are several cases where the use of large character fields is problematic, mostly in terms of a lack of query support. Refer to each of the supported database sections for details on the column type used for storing the values in these attributes.

### Isolation level

All of the enterprise beans accessed in a transaction must specify the same isolation level. If your application contains enterprise beans that are used in the same transaction as the rules, you must do one of the following:

- Change the BRBeans enterprise beans (`Rule`, `RuleFolder`, and `RuleHelper`) to the same isolation level as your beans.
- Change your beans to the same isolation level as the BRBeans enterprise beans.
- Place the BRBeans enterprise beans in a different database than your enterprise beans and configure the application to run using the two-phase commit protocol. This causes the beans to run in different transactions; thereby removing the restriction that they need to have the same isolation level.

## Oracle considerations

### Large character data

The preferred Oracle data type for storing large character objects is `CLOB`. However, Oracle does not allow a `CLOB` to be queried. Because of this, a data type of `VARCHAR2` is used by Business Rule Beans (BRBeans). A specific length must be specified when specifying `VARCHAR2`. The maximum length for a `VARCHAR2` is 4000 bytes.

To determine the default size of VARCHAR2, look in the table.ddl file that was generated when you deployed the code. If the default size is not acceptable for your application, you can do one of the following:

#### **Increase the size of the columns**

Keep in mind that maximum size for a VARCHAR2 in Oracle is 4000. Increase the column size either by changing the value in the **create table statement** or by changing the schema mapping and deploying the BRBeans JAR file.

#### **Change the schema mapping to specify CLOB**

Do this for any of the attributes that you do not wish to query and then deploy the BRBeans JAR file.

#### **Isolation level**

The default isolation level is REPEATABLE\_READ. Oracle does not support this isolation level. Therefore, the IBM WebSphere runtime environment converts this to the next highest isolation level, which in this case is SERIALIZED. Be aware that this isolation level tends to be overly restrictive as it prevents two clients from reading data at the same time. The BRBeansOracle.jar file specifies an isolation level of READ\_COMMITTED.

### **Sybase considerations**

#### **Allowing null values**

By default, Sybase does not allow null values in string columns (like VARCHAR, TEXT, and so on). You can change this default value for a database using "isql" by issuing the following command:

```
sp_dboption databasename, "allow nulls by default", true
```

In this example "databasename" is your database name.

#### **Large character data**

The large character data fields are stored in a column of type TEXT. Sybase allows TEXT fields to be queried only using the SQL "LIKE" operator. Queries against the columns that perform the SQL "IS NULL" or "IS NOT NULL" operations are not allowed by Sybase. The alternative is to specify a column type of VARCHAR. However, the maximum allowed size for a VARCHAR in Sybase is 255 characters. This is not considered a large enough value for storing firingParameters, initParameters, descriptions, and so on.

If performing "IS NULL" and "IS NOT NULL" type queries is important and the 255 character limitation is acceptable, change these column types to VARCHAR. To accomplish this, alter the schema mapping for the Rule bean and then deploy the BRBeans JAR file.

The query APIs (in the com.ibm.websphere.brb.query package) allow for "IS NULL" and "IS NOT NULL" type queries to be performed on several of these fields. In addition, the Rule Management Application allows the firing parameters to be queried in this manner. These queries fail on Sybase with the default column type of TEXT.

### **Informix considerations**

#### **Large character data**

The preferred Informix data type for storing large character data is CLOB. However, Informix does not allow a CLOB to be queried. Because of this, a data type of LVARCHAR is used by BRBeans. The maximum length for an Informix LVARCHAR is 2,000 characters. If 2,000 characters is not acceptable and your application does not need to query these data types, you can change the schema mapping to specify CLOB. Then, deploy the BRBeans JAR file.

#### **Custom properties for the data source**



When configuring the data source for your application, you must specify the following properties:

- `ifxIFXHOST`=Name of the physical machine on which the Informix instance is installed
- `serverName`=Informix instance name
- `portNumber`=Port number for which the Informix instance is configured
- `informixLockModeWait`=500

A setting of 500 causes a connection to wait for up to 500 seconds for a lock. If you have a busy system, this wait can appear to be a system hang. This setting has the same effect as running SET LOCK MODE TO WAIT 500 on the connection.

**Note:** The previous configuration values are subject to change. Consult your Informix documentation for updates.

## Rule Management Application

The Rule Management Application (RMA) is a tool that assists the user in performing high-level administration of rules and rule folders.

This includes the capability to create, modify, delete, import, or export rules or rule folders. The RMA tool can be used initially by the programmer to define rules interactively, and then by the domain analyst for rule management tasks. The main window for the RMA is the Rules Browser.

The column on the left side of the Rule Browser window shows a nested hierarchy of all of the existing rule folders. Click one of these folders to display the rules it contains. The names of these rules appear in the right column.

Navigate as you would in a typical file-management browser.

1. Click the "+" icon to expand by one level and click the "-" icon to collapse it.
2. Click a rule or folder name to highlight it, right-click the rule or folder name to launch a list of actions, or select an option from the main menu.

RMA is designed to be a general purpose tool for interactive management of rules. Alternatively, you can write your own user interface that is tailored more specifically to the domain in which you work. For instance, a domain-specific user interface can provide more help to the user in the task of managing rules than a general purpose tool such as RMA. If you plan to write your own user interface, refer to the `RuleImplementor` of the BRBeans Javadoc for assistance.

## Rule management APIs

Business Rule Beans (BRBeans) provide a set of APIs to perform rule management tasks programmatically. These tasks include creating, deleting, and updating rules and folders. These APIs are provided to simplify the interaction with the BRBeans enterprise beans. Use these APIs to perform rule management tasks instead of coding directly in the EJB interfaces.

The rule management APIs consist of the classes in the `com.ibm.websphere.brb.mgmt` package. You might use the following main classes:

### IRule

This is the interface used to access the object representing a business rule in BRBeans. It provides methods to read and update attributes of the rule, to delete the rule, and to make a copy of the rule. The methods to create rules are in the `IRuleFolder` interface since you must always create a rule and specify a particular folder in which it will reside. In the rule management APIs, any time you get a rule, you have the option to receive a reference to the enterprise bean itself or to receive a local copy of the data contained in the enterprise beans. Regardless of which option you choose, the `IRule` interface can be used to access the returned object. If a local copy is requested, it is possible to cast the returned object to an `IRuleCopy`. `IRuleCopy` extends `IRule` and adds a couple additional methods to those defined by `IRule`. See the **IRuleCopy** section for more details.



## **IRuleCopy**

This is the interface used to access a local copy of the enterprise bean that represents a business rule. An object implementing this interface is returned from rule management API methods if you ask for a local copy of the rule. The main reason for requesting a local copy is performance. Calling a method on a local copy is much faster than calling the method on the actual enterprise bean. If you need to access several different rule attributes, this may make a big difference. Similarly, when updating a rule, all updates can be sent to the enterprise bean in one method call instead of many. The individual set methods are called on the copy and then the `updatePersistentRule()` method is called to actually send the updates to the enterprise bean.

## **IRuleFolder**

This is the interface used to access the object representing a rule folder. It provides methods to create, delete, and find rules and subfolders. It also provides methods to move and rename the folder, and to get the parent folder. The `IRuleFolder` representing the root folder is generally what you start with when performing rule management tasks. Once you have the root folder you can navigate up and down the folder hierarchy and access rules contained within the folders.

## **RuleMgmtHelper**

This is a helper class intended to contain methods that are of general use for performing rule management tasks. Currently, the only methods available are used to get the `IRuleFolder` representing the root folder. The root folder is normally the starting point for performing rule management tasks.

## **IParameter**

This is the interface used to represent an initialization or firing parameter stored in a Rule EJB. Every parameter has a user description and a value that are accessible from this interface. The following classes are provided to implement the `IParameter` interface:

### **ConstantParameter**

This is the most common type of parameter. It represents a single constant value that is to be passed as an initialization or firing parameter.

### **MethodCallParameter**

This class represents a parameter whose value is determined by calling a method on the target object. The method to call must be a public method and must take zero parameters. This is only used for firing parameters.

### **TriggerPointParameter**

This class represents a parameter that is retrieved from one of the trigger point firing parameters. This is mainly used for reordering the firing parameters passed on the trigger point. This is only used for firing parameters.

For more details on the rule management interfaces, including a number of coding examples, refer to the `com.ibm.websphere.brb.mgmt` package in the BRBeans Javadoc.

## **BRBeans performance enhancements**

Externalizing business logic using Business Rule Beans (BRBeans) has many benefits, but does not come completely without a cost. Since every business rule is represented by an enterprise bean, then, in general, every rule trigger is performed in three parts:

1. a query is performed to find the enterprise beans that represent the rules to be triggered.
2. a remote method call is performed on the EJB instance to trigger the rule.
3. a remote method call is made to determine whether to fire the rule locally or remotely.

The first two steps both require server processing so processing can become rather slow.

This section documents the following ways to improve performance:

- Performance enhancements through caching
- Performance enhancements using indexes
- Performance enhancements by changing the firing location

## Performance enhancements through caching

The Business Rule Beans (BRBeans) framework incorporates a cache on the client side; that is, wherever the `trigger()` method on the `TriggerPoint` object is called. This cache is scoped to the Java virtual machine (JVM) in which the client is running so that any trigger calls performed in a particular JVM use the same cache and two triggers performed in different JVMs use two different caches. The BRBeans cache caches the results of all of the queries performed to find a set of rules to be triggered. The next time a trigger is performed in that JVM with the same rules specified, the rules are found in the cache and the query does not require server processing.

Once the rules are found in the cache they are triggered, either locally or remotely, depending on how they were defined. If a rule found in the cache is specified to be triggered locally, then the entire trigger process for that rule is performed on the client without calling the server. Even if the rule is specified to be triggered remotely, finding the rule in the cache eliminates one call to the server since the query is not performed on the server.

The BRBeans cache can improve performance greatly, however it has one disadvantage: changes made to rules are not recognized immediately.

When a change is made to a rule on the server, there is no way to inform all of the potential clients that something has changed and that they may need to refresh their caches. Thus, the client cache must check periodically to see if anything in the persistent rule data has changed. This is implemented by associating a polling frequency with the cache. This polling frequency specifies an interval of time that the cache waits before checking to see if anything has changed. The next time a trigger is performed after a polling interval has passed, the cache checks to see if any changes have been made to the persistent rule data stored on the server. If no changes have been made, then the cache is not refreshed. If any changes have been made, the entire cache is cleared so that the changes are picked up. Thus, changes to the rules are only picked up by the cache after a polling interval has passed.

The default polling frequency is 10 minutes. The user can change this value by changing the single initialization parameter specified for the special rule named `com/ibm/websphere/brb/BRB CacheRule`. The value for this initialization parameter is in the following format: `hh:mm:ss`

hh stands for hours, mm stands for minutes, and ss stands for seconds.

Thus the default of 10 minutes is specified by a value of `00:10:00`. To specify a polling frequency of, for example, 1 hour, 30 minutes, specify `01:30:00`

When this value is changed, it does not take effect until the previous polling interval has passed. Thus, if the previous polling interval is set to 24 hours and the polling frequency is changed to 1 hour, the new frequency does not take effect until the previous 24 hour polling interval passes. The only other ways to get the new frequency to take effect are to either restart the client (since this causes the cache to be re-initialized from scratch) or have the client code call the `refreshCache()` method on the `TriggerPoint` object. If there is more than one client JVM performing triggers, this must be done for each client since each JVM has its own cache.

**Note:** There is only one `BRB CacheRule` and this rule applies to all clients. There is no way to set different polling frequencies for different clients.

Caching can be disabled for a particular `TriggerPoint` object using the `disableCaching()` method. After the `disableCaching()` method is called any triggers performed using that `TriggerPoint` object can not use the cache. Triggers performed using other `TriggerPoint` objects are not affected.

## Performance enhancements using indexes

Creating an index over the database table that is used to store rules is an important way to improve the performance of rule queries. It is recommended that an index be created over the rulename column of the table containing the rules. This greatly improves the performance of rule-triggered queries that are looking for a rule or rules with a specific name. The index saves the query from searching every row in the table. Refer to the documentation for your database for instructions on how to create an index.

## Performance enhancements by changing the firing location

The Business Rule Beans (BRBeans) framework allows you to specify where to fire a particular rule. This determines where the rule implementor is instantiated and invoked. The following lists the possible values for the firing location:

**Local** Fires the Java rule implementor in the same JVM in which the trigger was performed.

### Remote

Fires the Java rule implementor on the server where the rules exist.

### Anywhere

Tries to fire the Java rule implementor locally. If the Java rule implementor cannot be found, then it fires the the implementor remotely. This is the default value.

For simple rule implementors that do not perform any server-intensive work, specifying **Local** usually results in the best performance. This is true both without and with caching. A complete comparison of local firing versus remote firing must consider four cases: local and remote firing without caching and local and remote firing with caching. A description of these four cases follow:

### Remote call without caching

Without caching, the work done to fire a rule remotely involves the following:

1. Finding the rule.
2. Determining whether the rule is to be fired locally or remotely.
3. Calling fire on the remote rule.

Each of these three operations requires a remote call to the server.

### Local call without caching

Without caching the work done to fire a rule locally involves the following:

1. Finding the rule.
2. Determining whether the rule is to be fired locally or remotely.
3. Calling fire on a local copy of the rule.

This requires only two remote calls. Firing locally results in a savings of one remote method call.

### Local call with caching

With caching, local firing results in even more dramatic improvements. The work done to fire a rule remotely involves the following:

1. Finding the rule. This involves a search of the local cache and does not involve calling the database.
2. Calling fire on the remote copy of the rule found in the server's cache.

This requires only one remote method call.

### Remote call with caching

The work done to fire a rule locally with caching involves the following:

1. Finding the rule. This involves a search of the local cache and does not involve calling the database.
2. Calling fire on the local copy of the rule found in the cache.

This does not involve any remote method calls. The entire rule firing process takes place locally without remote method calls. To get the full benefit of the BRBeans cache, use local firing. However, remotely fired rules still benefit from the cache due to the elimination of the query on the server and the elimination of the remote call to determine whether the rule is being fired locally or remotely.

There may be some cases where a rule implementor must perform some work that requires significant interaction with the server. In these cases, it may be beneficial to have rules using this rule implementor defined to be fired remotely. This might make the server interaction performed by the implementor more efficient.

**Note:** In addition to performance, maintenance also must be considered in relation to specifying a firing location. The rule implementor classes for rules that are defined to be fired locally must be present on any client system that tries to fire those rules. Otherwise, the implementor cannot be instantiated when the rule is fired. This can result in maintenance problems when the rule implementors are changed since they must be updated on many different systems.

---

## Developing BRBeans

Although the development tasks in this article are shown in sequential order, the exact sequence is left to your discretion. In particular, you might choose to create the business rules before or after the trigger point is placed. Also, you can write the rule implementor before or after creating the actual business rules. However, if you do not have the rule implementor at the time that the business rule is created, then you cannot complete the rule implementor field or the initialization parameters in the rule. However, you can leave the business rule marked as not ready until you can complete that information. For this reason, we have chosen to list the task of writing the rule implementor first. Complete the following tasks to develop Business Rule Beans:

1. Determining where to place the trigger point Determine where there are points of variability and where business logic must be externalized as part of application analysis and design process.
2. Placing a trigger point in the application code Add code to the application to invoke the trigger point framework, to find business rules, and to fire the rules.
3. Administering strategy objects to control triggers Control the process of triggering the rules using a set of strategy objects.
4. Implementing the business rules Use a rule implementor, written in Java, to implement the BRBeans `RuleImplementor` interface. Also, create business rules invoked by the rule implementor.

It is possible to develop your own customized strategy objects and rule implementors. See Customized strategy objects and Customized rule implementors for more information.

## Determining where to place a trigger point

To determine where to place your trigger points, you can use either the case analysis or the Object Interaction Diagrams (OIDs) method. The following are examples of methods that you can use to determine where to place a trigger point:

### Use case analysis

Trigger points can be found during analysis by inspecting the use cases or user interaction scenarios that are typically developed as statements of requirement as input to the analysis process. A fragment of a use case is shown below:

*The vehicle is entered into the system or chosen. The customer service representative attempts to locate the named driver in the system. If the driver is not found, she or he is added to the system and then picked.*

*Otherwise the found driver is simply picked. If the vehicle is an auto, anyone between the ages of 16 and 75 can be picked as a driver. If the vehicle is a truck, only drivers 16 to 70 years old can be picked. And if the vehicle is a motorcycle, drivers 14 to 65 can be picked.*

*After the driver has been picked, a rate quote can be performed...*

To identify potential trigger locations in use case analyses such as this one, look for **keywords** such as:

- "if X is in a special category Y" (For example, "if the vehicle is a truck")
- "except when"
- "unless"
- "depends on"

## Object Interaction Diagrams (OIDs)

OIDs that are based on use cases can yield a number of observable patterns that can be used to identify trigger points fairly easy. The following are some of the rules to look for and where the trigger point might be placed:

- Validation of edits on create methods.
- Validation of edits on set methods.
- Referential integrity of edits on methods that set references.
- Cardinality checks at a consistency point (a point in time where all of the data is expected to be self consistent).
- Required fields checks at a consistency point.
- Cross field edits at a consistency point.
- Constraints or derivations that have a high potential for reuse (especially if the algorithm is complex) at any appropriate point.
- Constraints or derivations that a business desires to be consistent across applications (at any appropriate point).
- Constraints or derivations where the business wants to decouple the maintenance cycle for a rule from the maintenance cycle for the code (at any appropriate point).

By using either method, you will be able to identify where to locate trigger points to use Business Rule Beans (BRBeans) effectively.

Once the trigger points have been identified, place the trigger point or points. See [Placing a trigger point](#) for more information.

## Placing a trigger point in the application code

Before placing a trigger point, review the following topics:

- Trigger point framework
- Determining where to place a trigger point

The `TriggerPoint` class is the primary interface of the Business Rule Beans (BRBeans) Trigger Point framework. The class is used to transfer control to the Trigger Point framework to find and fire the rules specified in the application's trigger point.

1. Create an instance of the `com.ibm.websphere.brb.TriggerPoint` class. All rule triggers must be performed against an instance of this class. Also, set any desired strategies on the `TriggerPoint` instance.
2. Gather together the parameters to be passed on the trigger.

For the simple `trigger()` and `triggerClassifier()` methods this includes the following:

### **An optional target object**

This can be used to specify an object that is to be the target of the rule's algorithm. Whether this is needed depends completely on the design of the rule implementor being used.

### The firing parameters for this rule trigger

This is an array of run-time parameters needed by the rule to satisfy its business purpose. The exact set of required firing parameters is determined by the rule implementor that is used by the rule.

**Note:** Any firing parameters defined on the rule itself will override whatever is passed here.

### Information identifying the rule or rules to be triggered

Normally this is either a single String containing the name of the rule to be triggered or an array of Strings each element of which is the name of a rule to be triggered. However if a custom finding strategy is being used, this could be whatever information it needs in order to find the correct rules.

The `triggerSituational` method differs in that it takes two sets of firing parameters and two sets of rule identification information: one set for the classifier rules and one for the classified rules.

3. Invoke the `trigger()`, `triggerClassifier()`, or `triggerSituational()` method of the `TriggerPoint` instance. This will actually trigger the rule or rules.
4. Process the results of the triggered rule or rules.

Examples of how to code a trigger point call are provided in the following topics:

- Simple trigger point
- Classifier trigger point
- Situational trigger point

For a detailed description of the trigger point programming interfaces refer to the `Trigger Point` class in the BRBeans Javadoc.

## Administering strategy objects to control triggers

Strategy objects are used to alter `TriggerPoint` functions. The two simple trigger methods, `trigger()` and `triggerClassifier()`, perform their function in the following sequence:

1. Find the rules.
2. Filter out those rules which are not desired.
3. Fire the remaining rules.
4. Combine the results and return to the caller.

The complex trigger method `triggerSituational()` does this sequence of steps twice; the first sequence finds the classification to feed into the second sequence.

Default strategy objects already are defined for each of the four `TriggerPoint` steps and they are used if none are specified explicitly. For each of these steps, there are at least two strategy objects used, one for triggering classifier rules, and one for triggering non-classifier rules. For the filtering step, there are actually three pairs of strategies that are used, based on the number of rules which the finding strategy returns (zero, one, or multiple).

While the sheer number of strategies that are available can be intimidating (twelve different strategy classes can be set), very few will need updating. Most users will modify the filtering strategies or the combining strategies only.

A number of predefined strategy objects are provided and can be used for the majority of cases. Although the following strategies are described separately, they have a definite dependency on each other. For example, the `FilteringStrategy` filters rules from the `FindingStrategy`; the `FiringStrategy` uses the results of the `FindingStrategy` to operate; and the `CombiningStrategy` depends on the results of the `FiringStrategy`.

### FindingStrategy

The `FindingStrategy` accesses the data store and returns those rules that meet the search criteria specified. See the `FindingStrategy` method for more information.



### **FilteringStrategy**

The FilteringStrategy takes the list of rules that were found by the FindingStrategy and filters out the rules that should not be fired. See the FilteringStrategy method for more information.

### **FiringStrategy**

The FiringStrategy takes the rules that were found by the FindingStrategy, (possibly modified by the FilteringStrategy), fires them each in order, and returns an array containing the results of each rule. See the FiringStrategy method for more information.

### **CombiningStrategy**

The CombiningStrategy takes the results of the rules that are fired by the FiringStrategy and combines them to form a reasonable result to the TriggerPoint caller. See the CombiningStrategy method for more information.

The Java classes for these strategy objects are defined in the `com.ibm.websphere.brb.strategy` package. Static constants also are defined in the interfaces for the various strategies. This allows easy access to instances of the strategy classes to set them on the TriggerPoint.

Also, it is possible to write your own strategy class if the supplied ones do not perform the function you need. See Customized strategy objects for more details.

## **Implementing business rules**

After determining where to place a trigger point, placing the trigger point, and defining your strategy objects, you must provide a method to implement the business rules and then create the rules. In addition, you might choose to organize your rules by creating rule folders.

**Note:** Although the development tasks in this article are shown in sequential order, the exact sequence is left to your discretion. In particular, you might choose to write the rule implementor before or after creating the actual business rules.

The business rule encapsulates and externalizes the business logic for the rule and any data that parameterizes the rule. Complete the following process to implement business rules:

1. Provide an implementation, called a rule implementor, for each business rule that you create. The rule implementor provides the actual business logic for the rule, implemented in Java. The rule implementor's `fire()` method is called when the business rule is triggered to actually perform the processing for the rule. Several generic rule implementors are shipped with BRBeans, which might be useful in some situations. However, if these rule implementors do not meet your needs, you can write your own rule implementor. Refer to the section on writing your own rule implementor for details.

2. Use the Rule Management Application (RMA) to create the business rule

- a. In the **Rule Browser** window, select the folder where you want the new rule to be created.
- b. From the main menu, click **File > New > Rule**.

In the **New Rule** properties window, use the following tabs to define the rule:

#### **General**

Use this tab to enter general information about the rule.

#### **Implementation**

Use this tab to define the manner in which the rule is implemented.

#### **Description**

Use this tab to define the purpose and intent of the rule.

#### **Dependent Rules**

Use this tab to specify the rules that the newly created rule will depend upon.

**Other** Use this tab to establish precedence, and enter information that is relevant to you, but does not fit into any other category.

- c. To complete the creation of the rule, click **OK**. If there are any mandatory fields still undefined, either go back and give them a value, or make the rule unavailable for use (see **Status** in the **General** tab for more information on this).



3. Create a rule folder using the Rule Management Application to contain the new business rule
  - a. In the **Rule Browser** window, select the folder where you want the new folder to be nested.
  - b. From the main menu, click **File > New > Folder**. A new folder appears in the folder hierarchy in edit mode. Enter a folder name and hit the **Enter** key.

---

## Assembling applications for use with BRBeans

When you are ready to ship your application, include a Business Rule Beans (BRBeans) JAR file in your EAR file. There are several of these JAR files in the <WAS\_HOME>/BRBeans directory; one for each supported database. Each name reflects the database type that it uses (for example, BRBeansDB2.jar). These JAR files contain three enterprise beans with the following JNDI names:

- brbeans/application/Rule
- brbeans/application/RuleFolder
- brbeans/application/RuleHelper

In your EAR files, complete the following steps:

1. Change the JNDI names of the BRBeans enterprise beans to make them unique for your application. For example, if your application is called MyApp, you could change the first one to brbeans/MyApp/Rule or com/MyCompany/MyApp/Rule.
2. Define EJB references to these three enterprise beans in any module where a trigger...() method exists in one or more of its classes. You can do this using the WebSphere Studio Application Developer tool or the Application Assembly tool. The Name field should contain the following and correspond to the enterprise beans listed above:
  - ejb/com/ibm/ws/brb/Rule
  - ejb/com/ibm/ws/brb/RuleFolder
  - ejb/com/ibm/ws/brb/RuleHelper

**Note:** The JNDI name on the **Bindings** tab should be the same as the JNDI names that you gave earlier to the enterprise beans.

Since the BRBeans enterprise beans refer to each other, there are also EJB references defined within the BRBeans JAR file itself. Each of the three BRBeans enterprise beans has two EJB references defined. These six references also need to be updated with the JNDI names you assigned earlier.

3. If you are not using the BRBeansCloudscape.jar file, skip this step. In the BRBeansCloudscape.jar file each entity enterprise bean has a resource reference defined for the data source it will use. You must update the JNDI binding for this reference to specify the JNDI name of the actual data source you want to use for the BRBeans entity enterprise bean. The shipped BRBeansCloudscape.jar file contains a dummy JNDI name for each binding. There are two resource references that need to be updated: one for the Rule enterprise bean and one for the RuleFolder enterprise bean. You can use the WebSphere Studio Application Developer tool or the Assembly Toolkit to perform this update.

After you have an EAR file that includes your application code and the BRBeans enterprise beans, complete the following steps to install and run the application:

1. Deploy the EAR file to generate run-time code for the BRBeans enterprise beans and any other enterprise beans that you may have in your application. Deployment can be accomplished either as a separate step (using the ejbdeploy command) or as part of the process of installing the EAR file onto an application server. Refer to the base WebSphere documentation for more information about deploying an EAR file.
2. The deployment process generates a file named Table.ddl. Table.ddl contains database commands to create the database tables needed by the BRBeans entity enterprise beans. You must use your database software to create a database and then use the commands in the Table.ddl file to create tables in this database. Refer to your database software documentation for more information on the commands needed to set up a database and the database tables.

3. Create a JDBC provider and a data source to access the database created in the previous step. For all of the databases except IBM Cloudscape, the BRBeans enterprise beans are configured to require a Version 4 data source. For Cloudscape, you can use a normal data source.

**Note:** The JNDI name you specify for the data source also must be specified when you install the EAR file on the application server so that the server knows which data source to use. Refer to the base WebSphere documentation for more information on these topics.

4. Install the EAR file on an application server. To install the EAR file, either use command line tools or the WebSphere Administrative Console.

**Note:** The EAR file can be deployed either as a separate step or as part of the installation process on the server.

Refer to the base WebSphere documentation for more information on installing an EAR file on an application server.

5. Start the application using the WebSphere Administrative Console.

You now have an application installed and running using BRBeans.

---

## Managing rules

In Business Rule Beans (BRBeans), rule management involves making changes to the set of business rules used by applications. This can include any of the following activities:

- Implementing the business rules: Creating rules and rule folders
- Copying and moving rule or rule folders
- Working with Quick Copy
- Finding a rule
- Deleting rules
- Deleting rule folders
- Changing the properties of a rule
- Importing a rule
- Exporting a rule
- Renaming rules
- Renaming rule folders
- Formatting columns
- Changing the date and time format

There are two different interfaces that can be used for rule management:

### **Rule Management Application**

An external user interface that allows users to manage rules interactively. It provides a general purpose interface for managing rules where no assumptions are made about the content or implementation of the rules. For information on how to use the Rule Management Application, see [Starting the Rule Management Application](#).

### **Rule management APIs**

A programmatic interface that can be used by programmers to manage rules or to customize an external user interface. For more information on Rule management APIs, see the [com.ibm.websphere.brb.mgmt](#) package.

Rules can be managed in any way that makes sense for your application, but the BRBeans framework was designed with the following administrative paradigm in mind:

1. Understand the desired change in business behavior.
2. Inspect the application documentation (in particular information that indicates where trigger points are located) to understand where the changes need to be made in the system.
3. Inspect the corresponding set of existing business rules using the Rule Management Application (or your own custom management application, if you have one) to understand which rules need to change.

4. Use the Rule Management Application, on a test system, to create one or more new rules that implement the required new behavior. Give these rules the correct name so that they are triggered by the appropriate trigger point. Also, make sure that these new rules are currently in effect.
5. Withdraw (by setting the end date of the rule), on the test system, all of the rules that are to be superseded.
6. Test the application to ensure that it behaves as expected.
7. Export the new rules using the Rule Exporter on the test system. Schedule the rules to become effective at the correct point in time.
8. Export the rules to be superseded using the Rule Exporter on the test system. Set them to expire when the new rules come into effect.
9. Import the new rules using the Rule Importer on the production system. This creates the new rules and schedules them to become effective at the date you specified when you exported them.
10. Import the rules to be superseded using the Rule Importer on the production system. This puts the new end date into the existing rules on the production system and sets them to expire on the specified date.

For more information on the overall topic of Externalizing business rules, see Using Business Rule Beans

## Starting the Rule Management Application

Review the overview topic Managing rules

To administer BRBeans, use the Rule Management Application (RMA). To launch the RMA, complete the following steps:

1. Open a command window and change to the following directory: `<WAS_HOME>/bin`
2. For Microsoft Windows platforms, enter `rulemgmt.bat <properties-file>`. For UNIX platforms, enter `rulemgmt.sh <properties-file>` where `<properties-file>` is a fully qualified name of a file containing port, host, and JNDI names used for the Business Rule Beans (BRBeans) enterprise beans. If you are using localhost port 2809 and you are using the default JNDI names for the BRBeans enterprise beans, you can specify the following in the default properties file:  
`<WAS_HOME>/bin/brbeansDefaultProperties`. For a full definition of the contents of this file, see the BRBeans properties file.

## Copying or moving rules or rule folders

Copy or move rules or rule folders either by cutting and pasting, or dragging and dropping.

To copy or move rules or rule folders, proceed as follows:

### Cutting and pasting

Use the menu commands (**Edit > Copy**, **Edit > Cut** and **Edit > Paste**) or keyboard commands (**CTRL-C**, **CTRL-V**, and **CTRL-X**).

### Dragging and dropping

Highlight the rule or rule folder you want to copy. Then press and hold the right mouse button, drag the cursor to the target location, and release. Select **Copy** or **Move** from the list.

**Note:** A rule also can be copied so that the copy replaces the existing rule at a specified date. This is referred to as a Quick Copy.

See Managing rules for other tasks related to the management of your rules.

## Working with Quick Copy

Use **Quick Copy** to make a copy of a rule that will replace the existing one on a specified date.

For example, suppose that we have an "isSeniorCitizen" rule. Currently a person is considered a senior citizen if they are 62 years of age or older. Starting on January 1, 2002, we are going to change this to 65. Use Quick Copy to specify the new date, and change the age from 62 to 65. The current rule will be set to expire on the same date that the new rule will take effect with the new senior citizen age defined as 65.

To use **Quick Copy** from the Rule Browser or Search Results window, proceed as follows:

1. Select the rules you want to **Quick Copy**.
2. From the main menu, click **Edit > Quick Copy**.
3. In the Quick Copy window, specify in the following fields how the copy will differ from the original:
  - Start Date For New Rule**  
Enter the date that the new rule will replace the existing rule. Use the date and time format that is shown.
  - Change parameter values for new rule**  
Enter the new parameter values.
4. Click **OK** to finish.

See Managing rules for other tasks related to the management of your rules.

## Finding a rule

To search for a specific rule using the Rule Management Application **Find** function, proceed as follows:

1. To search through all rules in all folders:
  - a. From the main menu of the Rule Browser, click **Edit > Find**.
  - b. Use the tabs in the Find Rules window to determine your search criteria.
2. To search a specific folder:
  - a. Right-click the folder and select **Find** from the list.
  - b. Use the tabs in the Find Rules window to determine your search criteria.

The results of your search are displayed in a Search Results window.

See Managing rules for other tasks related to the management of your rules.

## Deleting rules

To delete rules from the Rule Browser or Search Results window, proceed as follows:

1. Select the rules you want to delete.
2. From the main menu, click **File > Delete**.
3. Click **Delete** and then confirm the delete request.

**Note:** You cannot delete `com/ibm/websphere/brb/BRB CacheRule` as this rule is needed by the Business Rule Beans run-time environment.

See Managing rules for other tasks related to the management of your rules.

## Deleting rule folders

To delete rule folders from the Rule Browser window, proceed as follows:

1. Select the folder you want to delete.
2. From the main menu, click **File > Delete**.
3. Click **Delete** and then confirm the delete request.

**Note:** You cannot delete the root folder or any of the folders in the path `com/ibm/websphere/brb`.

See Managing rules for other tasks related to the management of your rules.

## Changing the properties of a rule

To change the properties of a rule, perform the following steps in either the Rule Browser or Search Result window:

1. Highlight the rule you wish to edit.
2. From the main menu, click **File > Properties**.

In the Rule Properties properties window, use the following tabs to edit the rule's definition:

### **General**

Use this tab to edit general information about the rule.

### **Implementation**

Use this tab to edit the manner in which the rule is implemented.

### **Description**

Use this tab to edit the purpose and intent of the rule.

### **Dependent Rules**

Use this tab to edit the list of rules that the newly created rule will depend upon.

**Other** Use this tab to establish precedence, and enter information that is relevant to you, but does not fit into any other category.

3. To complete the editing of the rule, click **OK**. If there are any mandatory fields still undefined, either go back and give them a value, or make the rule unavailable for use (see **Status** in the **General** tab for more information on this).

See Managing rules for other tasks related to the management of your rules.

## Importing a rule

To import rules from an XML format, use the Rule Browser window and proceed as follows:

1. In the main menu, click **File > Import**.
2. In the Import Rules window, specify the file you want to import.
3. Click **OK**. Rules and rule folders are created as specified within the XML.

See Managing rules for other tasks related to the management of your rules.

## Exporting a rule

To export rules, use the Rule Browser or Search Results window and proceed as follows:

1. In the main menu click **File > Export**.
2. In the **Export Rules Wizard**, proceed as follows:
  - a. In the Specify Rules to Export window, select the rule or rules that you want to export and click **Next**.
  - b. In the Change Effective Dates On Exported Rules window, alter the start and end dates of the rule (if desired) and click **Next**.
  - c. In the Select File For Rule Export window, choose a name and location for the exported rule.
3. Click **Export** to finish.

See Managing rules for other tasks related to the management of your rules.

## Renaming rules

To rename rules, use the Rule Browser or Search Results window and proceed as follows:

1. Highlight the rule you want to rename.
2. From the main menu, click **File > Rename**.
3. Type a new name and press the **Enter** key. To cancel the name change while it is still in progress, press the **Esc** key.

See Managing rules for other tasks related to the management of your rules.

## Renaming rule folders

To rename rule folders, use the Rule Browser or Search Results window and proceed as follows:

1. Place the folder name in edit mode by performing one of the following tasks:
  - a. Right-click the folder and select **Rename** from the list.
  - b. Highlight the folder and click **File >Rename** in the main menu.
  - c. Click the folder name twice with a slight pause between each click.
2. Type a new name and press the **Enter** key. To cancel the name change while it is still in progress, press the **Esc** key.

**Note:** You cannot change the name of the root folder.

See Managing rules for other tasks related to the management of your rules.

## Specifying columns

To choose which columns you want shown in your Rule Browser window, perform the following steps in either the Rule Browser or Search Results window:

1. From the main menu, click **View > Specify Columns**.
2. In the Specify Column window, proceed as follows:

### To add a new column

Select one or more entries in the **Available columns** list and click the **Add** button. The selected entries are added to the end of the Columns displayed list.

### To remove a column

Select one or more entries in the Columns displayed list and click the **Remove** button. The selected entries are added to the end of the Available columns list.

### To reorder columns

Select one or more entries in the Columns displayed list. To move the entries towards the top of the list, click the Up arrow. To move them towards the bottom, click the Down arrow.

See Managing rules for other tasks related to the management of your rules.

## Changing the date and time format

To change the date and time format, use either the **Rule Browser** or **Search Results** window and proceed as follows:

1. In the main menu, click **View > Specify Date/Time Format**.
2. In the **Specify Date/Time Format** window, choose one of the following radio button options:

### Use default format for this locale

Use this option to adjust the date and time format to match the default setting of your current locale.

### Select a predefined format for the date and time

Use this option to choose from one of several existing date and time formats.

### Specify a custom format for the date and time

Use this option to determine your own format for your date and time display.

3. When the example in the lower left of the window meets your needs, click **OK** to finish.

See Managing rules for other tasks related to the management of your rules.

---

## Rule Browser

The Rule Browser is the main window of the Rule Management Application (RMA), which is the tool used to administer business rules for Business Rules Beans. The RMA is a simple graphic user interface that assists the user in the high-level administration of rules and rule folders. The column on the left side of the Rule Browser window shows a nested hierarchy of all of the existing rule folders. Open one of these rule folders to display the rules. The names of these folders appear in the right column.

To navigate through the information, perform the following actions:

- Click the **+** icon to expand the folder by one level; click the **-** icon to collapse it.
- Click a file name to highlight it, right-click it to launch a list of actions, or select an option from the main menu.

## File menu

This article describes the options available in the File menu window.

### New

Use the **New** option to create one of the following objects:

#### Folder

The **Folder** selection creates a new folder within the folder currently selected in the browser. To create a new folder, complete the following steps:

1. Select the folder, in the Rule Browser window in which to nest the new folder.
2. Click **File > New > Folder** from the main menu. A new folder appears in the folder hierarchy in edit mode.
3. Enter a folder name and press **Enter**.

**Rule** The **Rule** selection creates a new rule within the folder that is currently selected in the browser. To create a new rule, complete the following steps:

1. Select the folder, in the Rule Browser window, where you want to create the new rule.
2. Click **File > New > Rule** from the main menu.
3. Use the following options, in the New Rule Properties window, to define the rule. For more information on each of these tabs, see the associated help file.

#### General

Use the **General** tab to enter general information about the rule. For more information, click the **New Rule properties window: General tab** link under Related reference.

#### Implementation

Use the **Implementation** tab to define the manner in which the rule is implemented. For more information, click the **New Rule properties window: Implementation tab** link under Related reference.

#### Description

Use the **Description** tab to define the purpose and intent of the rule. For more information, click the **New Rule properties window: Description tab** link under Related reference.

#### Dependent Rules

Use the **Dependent Rules** tab to specify the rules that the newly created rule will depend upon. For more information, click the **New Rule properties window: Dependent Rules tab** link under Related reference.

**Other** Use the **Other** tab to establish precedence and enter information that is relevant to



you, but does not fit into any other category. For more information, click the **New Rule properties window: Other tab** link under Related reference.

**Note:** When you change the properties of a rule and there are undefined mandatory fields, either give them a value or make the rule unavailable for use. See "status" on the General tab for more information.

4. Click **OK** to complete the creation of the rule.

### Rule Browser Window

This selection opens a new Rule Browser window on your desktop and shows the content of the currently selected folder.

## Import

Use the **Import** option to import rules that are defined in a file and written in XML. The rules are imported into folders as specified within the XML.

### Importing rules

You can use the Rule Browser window to import from an XML format. To import rules, complete the following steps:

1. Click **File > Import** in the main menu.
2. Specify the file, in the Import Rules window, that you want to import.
3. Click **OK**. The rules are imported as specified within the XML.

For more information, see the **Import Rules window** link under Related reference.

## Export

Use the Export option to export a file in an XML format.

### Exporting rules

You can export rules from the Rule Browser or Search Results windows. To export rules, complete the following steps:

1. Click **File > Export**, in the main menu.
2. Proceed as follows using the Export Rules Wizard:
  - In the Select Rules to Export window, select the rule or rules that you want to export and click **Next**.
  - In the Change Effective Dates On Exported Rules window, alter the start and end dates of the rule, if desired, and click **Next**.
  - In the Select File For Rule Export window, choose a name and location for the exported rule.
3. Click **Export** to finish.

For more information on this Export Rules Wizard, see the **Select Rules to Export window** link under Related reference.

## Delete

Use the **Delete** option to delete the selected rule or rule folder. If a rule folder is selected, all of the rules and subfolders the folder contains also are deleted.

### Deleting rules

You can delete rules from the Rule Browser or Search Results windows. To delete rules, complete the following steps:

1. Select the rules that you want to delete using the Rule Browser window.
2. Click **File > Delete** from the main menu.
3. Click **Delete** and then confirm the delete request.

**Note:** You cannot delete `com/ibm/websphere/brb/BRB CacheRule` as this rule is needed by the Business Rule Beans run-time environment.

### Deleting a folder

To delete a folder, complete the following steps:

1. Select the folder you want to delete using the Rule Browser window.
2. Click **File > Delete** in the main menu.
3. Click **Delete** and then confirm the delete request.

**Note:** You cannot delete the root folder.

### Rename

Use the **Rename** option to rename the selected rule or rule folder.

### Renaming a rule

You can rename rules from the Rule Browser or Search Results windows. To rename a folder, complete the following steps:

1. Highlight the rule you want to rename.
2. Click **File >Rename** from the main menu.
3. Type a new name and press **Enter**. To cancel the name change while it is still in progress, press **Esc**.

### Renaming a folder

1. Place the folder name in edit mode by doing any of the following:
  - Right-click the folder and select **Rename** from the list.
  - Highlight the folder and click **File >Rename** in the main menu.
  - Click the folder name twice with a slight pause between each click.
2. Type a new name and press **Enter**. To cancel the name change while it is still in progress, press **Esc**.

**Note:** You cannot rename the root folder.

## Properties

Use the **Properties** option to modify the properties of the selected rule or rule folder and then click **OK**.

Use the following tabs, in the **Rule Properties** window, to define the rule. For more information on each of these tabs, see the associated help file.

### General

Use the **General** tab to enter general information about the rule. For more information, click the **Rule properties window: General tab** link under Related reference.

### Implementation

Use the **Implementation** tab to define the manner in which the rule is implemented. For more information, click the **Rule properties window: Implementation tab** link under Related reference.

### Description

Use the **Description** tab to define the purpose and intent of the rule. For more information, click the **Rule properties window: Description tab** link under Related reference.

### Dependent Rules

Use the **Dependent Rules** tab to specify the rules that the newly created rule will depend upon. For more information, click the **Rule properties window: Dependent Rules tab** link under Related reference.

**Other** Use the **Other** tab to establish precedence and enter information that is relevant to you, but does not fit into any other category. For more information, click the **Rule properties window: Other tab** link under Related reference.

**Note:** When you change the properties of a rule and there are undefined mandatory fields, either give the fields a value or make the rule unavailable for use. See "status" on the **Rule properties window: General tab** for more information.

## Validate

Use the **Validate** option to verify the correctness of the selected rules or folder. A report is generated and displayed in system browser window. Use the browser functionality to view, save, and print the results. The following items are validated:

### Start and End dates

Specifies a valid start date. The end date must be empty (never expires) or be later than the start date.

### Java rule implementor name

Specifies a valid and existing Java rule implementor name. The specified class must implement the RuleImplementor interface.

### Dependent rules

Verifies the existence of each specified dependent rule.

When this option is selected, the Validate Rules window displays. You must specify a Java class path to validate rule implementor classes.

## Close

Use the **Close** option to terminate the application and close the window.

## New Rule properties window: General tab

Use the **New Rule properties window: General tab** to enter general information about the rule. The following fields and options are available on this tab:

### Name and location

#### Folder name

(Required) Use the **Folder name** field to identify the folder in which to create the rule. To browse the existing folders, click the **ellipses** button to the right of the text field.

**Name** (Required) Use the **Name** field to give the rule a name. The name cannot include the forward slash '/' and must contain at least one non-blank character. The name cannot exceed the maximum length of the rule name column in the database table.

### Period when rule is in effect

#### Start date

(Required) Use the **Start date** field to specify the date and time that the rule will go into effect. If you do not specify the time, a value of midnight is used.

#### End date

Use the **End date** field to determine the date when this rule expires. If you do not specify a value, the rule never expires.

## Classification

Use this section to specify your rule's classification status. Choose one of the following options:

- Rule is not classified and does not perform a classification (default)
- Rule performs a classification
- Rule is classified with the following classification

## Status

### Rule is available for use

Select **Rule is available for use** when the rule is available for use by the Business Rule Beans run-time environment. This feature is useful when you have not finished creating the rule, but need to save your changes.

## New Rule properties window: Implementation tab

Use the **New Rule properties window: Implementation** tab to define the manner in which the rule is implemented.

**Note:** The rule contains the data, but it does not perform the implementation. Rather, the Java rule implementor implements the rule.

For example, suppose you want to create a rule that determines whether a given person is a senior citizen, 62 years old or older. To implement this rule, give the "com.ibm.websphere.brb.implementor.RuleGreaterThanEqual" Java rule implementor the value 62 and specify it as an initialization parameter. When the application fires the rule, the person's age is passed to the Java rule implementor as a firing parameter and 62 is passed as an initialization parameter. The person's age is compared against the initialization parameter of 62 and a value of true or false is returned from the Java rule implementor to the application. To change the age at which a person is considered a senior citizen, change the value of the initialization parameter.

The following fields and options are available on this tab:

**Java rule implementor:** (Required) Use the **Java rule implementor** field to specify a class to implement this rule. The initialization and firing parameters that are required are determined by looking at the documentation for the specified Java rule implementor.

**Firing location:** Use the **Firing location** field to determine where the rule is fired. You can fire the Java rule implementor on the server where the rules exist or fire it locally on the client machine. The client can be a servlet running on the server. Specify one of the following values for the firing location:

**Local** Use the **Local** option to fire the Java rule implementor local to the application that fired the rule.

### **Remote**

Use the **Remote** option to fire the Java rule implementor on the server where the rules exist.

### **Anywhere (default)**

Use the **Anywhere** option to attempt to fire the Java rule implementor locally first. If the Java rule implementor cannot be found, then it is fired remotely.

To choose the value of the firing location, you must take both performance and maintenance into consideration. Most rules perform better if they are run on the same Java virtual machine (JVM) as the application (locally). However, there might be cases where a Java rule implementor performs server-intensive tasks, in which case the rules might run better when they run on the server. To run locally, you must have all of the Java rule implementors installed locally. They must be accessible by the application that fires the rules.

**Initialization parameters:** The initialization parameters contain constant values passed to the rule implementor when it is initialized. Typically, the initialization parameters contain values that might change as your business practices evolve, such as the age at which a person is considered a senior citizen or the current interest rate for a loan.

To add a new initialization parameter to the list, click **Add** and fill in the fields in the Add Initialization Parameter window.

To edit an existing initialization parameter, highlight it, click **Change**, and fill in the fields in the Change Initialization Parameter window.

To delete an initialization parameter, highlight it and click **Delete**.

To change the order of the initialization parameters, highlight one and click the up or down arrows to move it to a new location.

**Firing parameters:** The firing parameters contain values passed from the trigger point in the application to the Java rule implementor at run time. You can alter the parameters coming from the application before passing them to the Java rule implementor. Typically, these parameters are left unaltered.

For example, when implementing the "isSeniorCitizen" rule that determines whether a person is a senior citizen, you might want to pass a person from the application to the rule as the target object. However, the "isSeniorCitizen" rule uses the RuleGreaterThanEqual Java rule implementor, which requires that you pass an integer value. You can alter the firing parameters to specify that the method "getAge" is called on the person object and pass that result to the Java rule implementor.

You must choose one of the following options:

**Pass firing parameters from trigger point unchanged**

The parameters specified in the trigger point of the application are passed to the Java rule implementor without being altered. This is the default value.

**Specify firing parameters**

The values specified in the table are passed to the Java rule implementor.

To add a new firing parameter to the list, click **Add** and fill in the fields in the Add Firing Parameter window.

To edit an existing firing parameter, highlight the firing parameter, click **Change**, and fill in the fields in the Change Firing Parameter window.

To delete a parameter, highlight the firing parameter and click **Delete**.

To change the order of the firing parameters, highlight one and click the up or down arrows to move it to a new location.

**Add Initialization Parameter window:** Use the Add Initialization Parameter window to add initialization parameters to a rule. The initialization parameters contain constant values passed to the Java rule implementor when it is initialized. Typically, the initialization parameters contain values that will change as your business practices evolve. These values might be the age at which a person is considered a senior citizen or the current interest rate for a loan.

To add an initialization parameter, proceed as follows:

1. Fill in the fields as needed.
2. Click **Add**. The initialization parameter is added and the window remains open to specify additional parameters.
3. When you are finished specifying initialization parameters, click **Close**.

The following fields and options are available in this window:

*Description:* Use the **Description** field to specify a description of the initialization parameter. This field can contain any information necessary to describe the purpose of the initialization parameter.

*Type:* Use the **Type** field to specify the type of data that is contained within this initialization parameter. The data itself is stored in the **Value** field (see the following field description). For example, if this initialization parameter specifies the age at which a person is considered a senior citizen, then the Type likely is an "Integer". If the initialization parameter specifies a company name, such as "IBM", then the Type is a "String".

The following values are available:

- String
- Character
- Byte
- Short

- Integer
- Long
- Float
- Double
- Boolean
- java.math.BigDecimal
- java.math.BigInteger
- Null Value

*Value:* Use the **Value** field to add a value for the parameter. For example, if the initialization parameter is intended to specify the age at which a person is considered a senior citizen, then this field might be set to 62.

**Change Initialization Parameter window:** Use the Change Initialization Parameter window to edit an existing initialization parameter.

The initialization parameters contain constant values passed to the Java rule implementor when it is initialized. Typically, the initialization parameters contain values that change as your business practices evolve. These values might be the age at which a person is considered a senior citizen or the current interest rate for a loan.

To edit an existing initialization parameter, modify the fields and click **OK**.

The following fields are available in this window:

*Description:* Use the **Description** field to specify a description of the initialization parameter. This field can contain any information necessary to describe the purpose of the initialization parameter.

*Type:* Use the **Type** field to specify the type of data that is contained within this initialization parameter. The data itself is stored in the **Value** field (see the following field description). For example, if this initialization parameter specifies the age at which a person is considered a senior citizen, then the Type likely is an "Integer". If the initialization parameter specifies a company name, such as "IBM", then the Type is a "String".

The following values are available:

- String
- Character
- Byte
- Short
- Integer
- Long
- Float
- Double
- Boolean
- java.math.BigDecimal
- java.math.BigInteger
- Null Value

*Value:* Use the **Value** field to add a value for the parameter. For example, if the initialization parameter is intended to specify the age at which a person is considered a senior citizen, then this field might be set to 62.

### **New Rule properties window: Description tab**

Use the **New Rule properties window: Description** tab to define the purpose and intent of the rule. All of the fields in this panel are optional and none of them are used by the Business Rule Beans run time. The following fields are available on this tab:

**Business Intent:** Use the **Business Intent** field to describe the business intent of this rule.

**Description:** Use the **Description** field to define a general description of the rule and its purpose.

**Original requirement:** Use the **Original requirement** field to compose a description of the original requirement that created this rule.

### **New Rule properties window: Dependent Rules tab**

Use the **New Rule properties window: Dependent Rules** tab to specify the rules that the newly created rule will depend upon.

To add names to the **Dependent rule names** field, proceed as follows:

1. Locate the dependent rule. You can do this in one of two ways:

#### **Browse for a rule**

If you are familiar with the location of the dependent rule, then click **Browse**. Navigate to the rule's location and highlight it.

#### **Find a rule**

If you are unfamiliar with the location of the rule, then click **Find**. This launches a Find Rules window in which you can specify options and then initiate a search. When you have located the rule, highlight it.

2. Click **Add**.

To change the name of a rule in the **Dependent rule names** field, highlight one or more rules and click **Edit**. Specify the new name for the rule and click **OK**.

To change the folder name of a rule in the **Dependent rule names** field, highlight one or more rules and click **Change Folder**. Fill-in the fields in the Change Folder window. All selected rules are updated with the newly specified values.

To delete a rule from the **Dependent rule names** field, highlight the rule and click **Delete**.

To change the order of the rules in the **Dependent rule names** field, highlight a rule and click the up or down arrows to move it to a new location.

### **New Rule properties window: Other tab**

Use the **New Rule properties window: Other tab** to establish precedence and enter information that is relevant to you, but does not fit into any other category. The following fields are available on this tab:

**Precedence:** Use the **Precedence** field to specify the relative priority when firing the rule. This value is used to order the rules from lowest to highest.

**User defined data:** Use the **User defined data** field to enter any additional text that you want to store. The Business Rule Beans run time does not use this field.

### **Import Rules window**

Use the Import Rules window to select and import a rule expressed in an XML format. The following fields and options are available in this window:

**File Name:** Use the **File Name** field to specify the name of a file that contains the rules that you want to import. To search for a file, click the **ellipses** icon to the right of the text entry field.

**Show output from rule importer:** Select **Show output from rule importer** to display detailed information about the rules that you want to import.

**Show rules to be created but do not create them:** Select **Show rules to be created but do not create them** to validate the XML prior to committing to the rule's creation. The rule XML runs through the importer but is not created.



**Update existing rules with the same primary key:** Select **Update existing rules with the same primary key** to update the rules with the same primary key. If this check box is clear, the rule is not imported if one is found with this same primary key. An error message is shown and the transaction in which this rule was created is not committed. The primary key is an optional tag within the XML and there is a possibility that a rule already exists on the system with this same primary key.

**Transaction Option:** The following is a list of transaction options:

**One transaction per rule**

Use the **One transaction per rule** option to start a transaction for each rule. If one rule fails to import, it does not prevent other rules in the specified file from being imported.

**One transaction for all rules**

Use the **One transaction for all rules** option to stop all of the rules from being imported if any rule fails to import successfully. Use this feature to ensure that all of your rules are in a consistent state. Typically, it is undesirable to have only a portion of the rules imported successfully since rules might have dependencies on other rules.

## Select Rules To Export window

This is the first of three windows in the Export Rules Wizard. Use the Select Rules To Export window to select the rules to export. After entering the appropriate information in the following fields, click **Next**.

**Note:** If you specify a folder, the tool exports the entire contents of the folder including its subfolders.

The following options are available in this window:

**Add:** Use the **Add** option to open a window in which you can type the fully qualified name of a rule or a rule folder. If there are multiple rules with the specified name, they are all exported.

**Find:** Use the **Find** option to open a Find Rules window in which you can specify your search criteria. If there are multiple rules with the same name, only the selected rules are exported.

**Browse button:** Use the **Browse** button to open a window and browse for rules to add to the list. If there are multiple rules with the same name, only the selected rules are exported.

**Remove button:** Use the **Remove** button to remove the selected rules from the list.

**Show output from rule exporter:** Select **Show output from rule exporter** to open a window that contains details about the export operation of the select rules. This window is shown when the export operation begins.

**Change Effective Dates On Exported Rules window:** This is the second of three windows in the Export Rules Wizard. Use the Change Effective Dates on Exported Rules window to alter the rule's start and end times. This procedure is useful when the application and the rules are tested on a development system prior to being deployed on a production system. You can change the dates of the rules and test on the development system using the current date, even if you plan to use the rules on the production system at a future date. The rules on the development system are not changed; only the exported version of the rules are changed. After determining whether to select the check boxes in the following descriptions, click **Next**.

The following fields and options are available in this window:

**Change start date and time on exported rules:** Select **Change start date and time on exported rules** to alter the start date of the rules to export. You must specify a valid date and time using the format shown. For example, the format might be "m/d/yy h:mm a" resulting in 10/23/01 1:45 PM.

**Change end date and time on exported rules:** Select **Change end date and time on exported rules** to alter the end date of the rules to export. You must specify a valid date and time using the format shown. For example, the format might be "m/d/yy h:mm a" resulting in 10/23/01 1:45 PM.

**Select File For Rule Export window:** This is the third of three windows in the Export Rules Wizard. Use the Select File For Rule Export window to choose a name and location for the exported rule.

1. Browse to an existing directory or create a new one.
2. Type the name of the file. Typically, the file will end with an ".xml" extension.
3. Click **Export**.

### Rule properties window: General tab

Use the **Rule properties window: General** tab to enter general information about the rule. The following fields and options are available in this window:

#### Name and location

##### Folder name

(Required) Use the **Folder name** field to identify the folder in which to create the rule. To browse the existing folders, click the **ellipses** button to the right of the text field.

**Name** (Required) Use the **Name** field to give the rule a name. The name cannot include the forward slash '/' and must contain at least one non-blank character. The name cannot exceed the maximum length of the rule name column in the database table.

#### Period when rule is in effect

##### Start date

(Required) Use the **Start date** field to specify the date and time that the rule goes into effect. If you do not specify the time, a value of midnight is used.

##### End date

Use the **End date** field to determine the date when this rule expires. If you do not specify a value, the rule never expires.

#### Classification

Use the **Classification** section to specify your rule's classification status. Choose one of the following options:

- Rule is not classified and does not perform a classification (default)
- Rule performs a classification
- Rule is classified with the following classification

#### Status

##### Rule is available for use

Select **Rule is available for use** when the rule is available for use by the Business Rule Beans run-time environment. This feature is useful when you have not finished creating the rule, but you want to save your changes.

### Rule properties window: Implementation tab

Use the **Rule properties window: Implementation** tab to define the manner in which the rule is implemented.

**Note:** The rule contains the data and it does not perform the implementation. Rather, the Java rule implementor implements the rule

For example, suppose you want to create a rule that determines whether a given person is a senior (62 years old or older). To implement this rule, give the "com.ibm.websphere.brb.implementor.RuleGreaterThanEqual" Java rule implementor the value 62 and specify it as an initialization parameter. When the application fires the rule, the person's age is passed to the Java rule implementor as a firing parameter and 62 is passed as an initialization parameter. The person's age is compared against the initialization parameter of 62 and a value of true or false is returned from the Java rule implementor to the application. To change the age at which a person is considered a senior citizen, change the value of the initialization parameter.

The following fields and options are available on this tab:

**Java rule implementor:** (Required) Use the **Java rule implementor** field to specify a class to implement this rule. The initialization and firing parameters that are required are determined by looking at the documentation for the specified Java rule implementor.

**Firing location:** Use the **Firing location** field to determine where the rule is fired. You can fire the Java rule implementor on the server where the rules exist or fire it locally on the client machine. The client can be a servlet running on the server. Specify one of the following values for the firing location:

**Local** Use the **Local** option to fire the Java rule implementor local to the application that fired the rule.

**Remote**

Use the **Remote** option to fire the Java rule implementor on the server where the rules exist.

**Anywhere (default)**

Use the **Anywhere** option to attempt to fire the Java rule implementor locally first. If the Java rule implementor cannot be found, then it is fired remotely.

To choose the value of the firing location, you must take both performance and maintenance into consideration. Most rules perform better if they are run on the same Java virtual machine (JVM) as the application (locally). However, there might be cases where a Java rule implementor performs server-intensive tasks, in which case the rules might run better when they run on the server. To run locally, you must have all of the Java rule implementors installed locally. They must be accessible by the application that fires the rules.

**Initialization parameters:** The initialization parameters contain constant values passed to the rule implementor when it is initialized. Typically, the initialization parameters contain values that might change as your business practices evolve, such as the age at which a person is considered a senior citizen or the current interest rate for a loan.

To add a new initialization parameter to the list, click **Add** and fill in the fields in the Add Initialization Parameter window.

To edit an existing initialization parameter, highlight it, click **Change**, and fill in the fields in the Change Initialization Parameter window.

To delete an initialization parameter, highlight it and click **Delete**.

To change the order of the initialization parameters, highlight one and click the up or down arrows to move it to a new location.

**Firing parameters:** The firing parameters contain values passed from the trigger point in the application to the Java rule implementor at run time. You can alter the parameters coming from the application before passing them to the Java rule implementor. Typically these parameters are left unaltered.

For example, when implementing the "isSeniorCitizen" rule that determines whether a person is a senior citizen, you might want to pass a person from the application to the rule as the target object. However, the "isSeniorCitizen" rule uses the RuleGreaterThanEqual Java rule implementor, which requires that you pass an integer value. You can alter the firing parameters to specify that the method "getAge" is called on the person object and pass that result to the Java rule implementor.

You must choose one of the following:

**Pass firing parameters from trigger point unchanged**

The parameters specified in the trigger point of the application are passed to the Java rule implementor without being altered. This is the default value.

**Specify firing parameters**

The values specified in the table are passed to the Java rule implementor.

To add a new firing parameter to the list, click **Add** and fill in the fields in the Add Firing Parameter window.

To edit an existing firing parameter, highlight the firing parameter, click **Change**, and fill in the fields in the Change Firing Parameter window.

To delete a parameter, highlight the firing parameter and click **Delete**.

To change the order of the firing parameters, highlight one and click the up or down arrows to move it to a new location.

**Add Firing Parameter window:** Use the Add Firing Parameter window to add a firing parameter to a rule. The firing parameters contain values passed from the trigger point in the application to the Java rule implementor at run time. You can alter the parameters coming from the application before passing them to the Java rule implementor. Typically, these parameters are left unaltered.

To add a firing parameter, enter a field description, select an appropriate option, and click **Add**. The parameter is added and the window remains open to specify additional parameters. When you finish specifying initialization parameters, click **Close**. The following fields and options are available in this window:

*Description:* Use the **Description** field to type a description of the firing parameter.

*Specify a type and value:* Use the **Specify a type and value** option to specify a constant value to pass to the Java rule implementor.

*Get value from method call:* Use the **Get value from method call** option to call the specified method on the target object.

*Get value from trigger point firing parameters:* Use the **Get value from trigger point firing parameters** option to get a specific value from the firing parameters that were specified in the trigger method in the application. An index to the original firing parameter must be specified. This index starts with 0; thus, specify the value 0 to pass the first firing parameter. Specify the trigger point parameter number.

**Change Firing Parameter window:** Use the Change Firing Parameter window to edit an existing firing parameter. The firing parameters contain values passed from the trigger point in the application to the Java rule implementor at run time. You can alter the parameters coming from the application before passing them to the Java rule implementor. Typically, these parameters are left unaltered.

To change a firing parameter, enter a field description, select an appropriate option, and click **Change**. The parameter is changed and the window remains open to specify additional parameters. When you finish specifying initialization parameters, click **Close**.

The following fields and options are available in this window:

*Description:* Use the **Description** field to type a description of the firing parameter. There are three types of firing parameters that can be specified.

*Specify a type and value:* Use the **Specify a type and value** option to specify a constant value to pass to the Java rule implementor.

*Get value from method call:* Use the **Get value from method call** option to call the specified method on the target object.

*Get value from trigger point firing parameters:* Use the **Get value from trigger point firing parameters** option to pass a specific value from the firing parameters that were specified on the trigger method in the application. An index to the original firing parameter must be specified. This index starts with 0; thus, specify the value 0 to pass the first firing parameter.

### **Rule properties window: Description tab**

Use the **Rule properties window: Description** tab to define the purpose and intent of the rule. All of the fields in this panel are optional and none are used by the Business Rule Beans run-time environment. The following fields are available on this tab:

**Business Intent:** Use the **Business Intent** field to describe the business intent of this rule.

**Description:** Use the **Description** field to define a general description of the rule and its purpose.

**Original requirement:** Use the **Original requirement** field to compose a description of the original requirement that created this rule.

### **Rule properties window: Dependent Rules tab**

Use the **Rule properties window: Dependent Rules** tab to specify the rules that the newly created rule will depend upon.

To add names to the **Dependent rule names** field, proceed as follows:

1. Locate the dependent rule. You can do this in one of two ways:

#### **Browse for a rule**

If you are familiar with the location of the dependent rule, then click **Browse**. Navigate to the rule's location and highlight it.

#### **Find a rule**

If you are unfamiliar with the location of the rule, then click **Find**. This launches a Find Rules window in which you can specify options and then initiate a search. When you have located the rule, highlight it.

2. Click **Add**.

To change the name of a rule in the **Dependent rule names** field, highlight one or more rules and click **Edit**. Specify the new name for the rule and click **OK**.

To change the folder name of a rule in the **Dependent rule names** field, highlight one or more rules and click **Change Folder**. Fill-in the fields in the Change Folder window. All selected rules are updated with the newly specified values.

To delete a rule from the **Dependent rule names** field, highlight the rule and click **Delete**.

To change the order of the rules in the **Dependent rule names** field, highlight a rule and click the up or down arrows to move it to a new location.

**Change Folder window:** Use the Change Folder window to change the folder name for all selected dependent rules. There are two options for changing the folder name:

*Specify a new folder name:* Use the **Name** field to specify the new folder name for all selected rules. All selected dependent rule names are updated to contain the new folder, exactly as specified.

*Replace common portion of the folder name:* Use this option to change the portion of the folder name that is common to all selected rules.

For example, suppose the following two rules exist:

- com/acme/insurance/health/ruleA
- com/acme/insurance/life/ruleB

You want to change the `com/acme/insurance` portion of the name to `com/acme/ins`. The `com/acme/insurance` portion will be changed while leaving the remaining portion of the rule name intact.

The **Common folder** field is set to the common portion of the folder name. In the preceding example, this field is set to `com/acme/insurance`. This value cannot be changed.

Specify the new folder name in the **Replace with** field. For the preceding example, specify `com/acme/ins`.

### Rule properties window: Other tab

Use the **Rule properties window: Other** tab to establish precedence and enter information that is relevant to you, but does not fit into any other category. The following fields are available in this window:

**Precedence:** Use the **Precedence** field to specify the relative priority when firing the rule. This value is used to order the rules from lowest to highest.

**User defined data:** Use the **User defined data** field to enter any additional text that you want to store. The Business Rule Beans run-time environment does not use this field.

### Validate Rules

Use the **Validate** option to verify the correctness of the selected rules or folder. A report is generated and displayed in system's browser window. You can use the browser functionality to view, save, and print the results. The following items are validated:

#### Start and End dates

A valid start date must be specified. The end date must be empty (never expires) or be later than the start date.

#### Java rule implementor name

A valid and existing Java rule implementor name must be specified. The specified class must implement the `RuleImplementor` interface.

#### Dependent rules

The existence of each specified dependent rule is verified.

You can specify a Java classpath to use to validate rule implementor classes for the rules. To add individual JAR files or directories containing classes to the **Classpath** field, click **Add** or **Browse**.

To specify the full name of a jar file or directory to be added to the classpath, click **Add**.

To search through the file system for a JAR file or directory to be added to the classpath, click **Browse**.

To change an existing classpath entry, highlight the entry and click **Edit**. Specify the new name for the classpath entry and click **OK**.

To remove a classpath entry from the list, highlight one or more entries and click **Remove**.

To change the order of the classpath entries, highlight one or more entries and click the Up or Down arrows to move it to a new location.

After specifying the classpath, click **OK** to generate the report.

### Edit menu

This menu describes the options available on the Edit menu window. The following options are available in this window:

#### Cut

Use the **Cut** option to move rules and rule folders.



## Copy

Use the **Copy** option to copy rules and rule folders. The following tasks can be accomplished using the **Copy** option:

- **Copying rules**

Copy or move a rule from one folder to another by either cutting and pasting it or dragging and dropping it.

- **Cutting and pasting**

- Use menu commands (**Edit > Copy**, **Edit > Cut** and **Edit > Paste**) or keyboard commands (**CTRL-C**, **CTRL-V** and **CTRL-X**) to copy or move a rule.

- **Dragging and dropping**

- Highlight the rule you want to copy. Then, press and hold the right mouse button, drag the cursor to the target location, and release. Select **Copy** or **Move** from the list.

**Note:** You also can copy a rule so that the copy replaces the existing rule at a specified date. This is referred to as a Quick Copy.

- **Copying rule folders**

Copy or move a rule folder and all its contents by either cutting and pasting it or dragging and dropping it.

- **Cutting and pasting**

- Use menu commands (**Edit > Copy**, **Edit > Cut** and **Edit > Paste**) or keyboard commands (**CTRL-C**, **CTRL-V** and **CTRL-X**) to copy or move a rule.

- **Dragging and dropping**

- Press and hold the right mouse button on the folder to be copied. Drag the cursor to the target location and release the mouse button. Select **Copy** or **Move** from the list.

## Paste

Use the **Paste** option to add cut or copied rules and rule folders.

## Find

Use the **File** option to search for a rule. A window opens in which you can specify your search criteria. If you would like to search a specific folder, right-click the folder and select **Find** from the list. For more information, click the **Find Rules window** link under Related reference.

- **Search the whole directory**

- To search the whole directory, complete the following steps:

1. Click **Edit > Find** from the main menu of the Rule Browser.
2. Determine your search criteria in the Find Rules window.

- **Search a specific folder**

- To search a specific folder, complete the following steps:

1. Right-click the folder and select **Find** from the list.
2. Determine your search criteria in the Find Rules window.

The results of your search are displayed in a Search Results window.

## Quick Copy

Use the **Quick Copy** option to make a copy of a rule that will replace the existing one on a specified date. You can modify the copy so that a new value goes into effect on the desired date and time while the old rule expires. For more information, see the **Quick Copy window** link in Related reference.

## Select All

Use the **Select All** option to facilitate rule selection.

## Deselect All

Use the **Deselect All** option to deselect rules.



## Quick Copy window

Use the Quick Copy window to make a copy of a rule that replaces the existing one on a specified date.

For example, suppose that you have an "isSeniorCitizen" rule. Currently, a person is considered a senior citizen if they are 62 years of age or older. Starting on January 1, 2002, you must change this to 65. Use Quick Copy to specify the new date and to change the age from 62 to 65. The current rule is set to expire on the same date that the new rule takes effect. The new senior citizen age is defined as 65.

**Note:** Use the **Quick Copy** function for simple changes only.

In the following fields, specify how the copy differs from the original:

**Start Date For New Rule:** Use the **Start Date For New Rule** field to enter the date that the new rule replaces the existing rule. Use the date and time format that is shown. For example, the format might be the following:

M/d/yy h:mm a.

**Change parameter values for new rule:** Use the **Change parameter values for new rule** field to add new parameter values.

## Automatically Update References window

Use the Automatically Update References window to indicate what you would like to do with rules that contain references to the rules being moved, copied, or renamed. You can choose whether to update references and whether to view the references before updating the rules. When the rule references are displayed, you can choose which references to update. The following options are available in this window:

**Automatically update references:** Select **Automatically update references** to indicate that references should be updated.

**Display the updates before changing the rules option:** Select **Display the updates before changing the rules option** to indicate whether to show the changes before updating the references. When you choose this option, the Update Dependent Rule References window is displayed, allowing you to select which references to update.

## Update Dependent Rule References window

Use the Update Dependent Rule References window to select the dependent rule references that must be updated as a result of the move, copy, or rename operation.

Each row in the table represents one dependent rule reference. The following information is shown:

### Rule Name

Specifies the name of the rule containing the dependent rule reference.

### Folder Name

Specifies the folder containing the rule specified by **Rule Name**.

### Dependent Rule Old Name

Specifies the name of the dependent rule that is to be changed as a result of this move, copy, or rename operation.

### Dependent Rule New Name

Specifies the new name to be given to the dependent rule as a result of this move, copy, or rename operation.

### Status

Specifies the status of the rule; in effect, scheduled, expired, or invalid.

### Start Date

Specifies the date and time when this rule will begin being used.

### End Date

Specifies the date and time of when this rule will expire (no longer be used).

The check box to the left of the rule name indicates whether the dependent rule reference will be updated. Select the check box to change the selection which indicates that the reference should be updated.

Use the **Select All** button to indicate that all dependent rule references should be updated.

Use the **Deselect All** button to clear the dependent rule references selection.

Click **Continue** after reviewing the list of dependent rule references and indicating which references to update. Click **Cancel** to cause the move, copy, or rename operation to be cancelled completely.

### References To Rules Being Deleted window

Use the References To Rules Being Deleted window to view dependent rule references when a rule or folder is deleted. Each row in the table represents one dependent rule reference. The following information is shown:

**Rule Name**

Specifies the name of the rule containing the dependent rule reference.

**Folder Name**

Specifies the name of the folder containing business rules.

**Dependent Rule Reference**

Specifies the name of the dependent rule that is being deleted.

**Start Date**

Specifies the date and time when this rule is implemented.

**End Date**

Specifies the date and time when this rule expires.

Click **Continue** to delete. Click **Cancel** to abort.

### Rule Names Already Exist window

The Rule Names Already Exist window displays when one or more rules are copied into a folder and existing rules are found with the same name. This window enables you to rename the rules before you proceed with the copy operation. The conflicting rule names do not need to be changed; you can create multiple rules with the same name.

Each line of the table represents one rule for which a name conflict has occurred. The table shows the following fields:

**Rule Name**

Specifies the name of the rule being copied which is the same as the name of a rule that already exists in the destination folder.

**Folder Name**

Specifies the folder containing the rule specified by Rule Name

**Status**

Specifies the status of the rule; in effect, scheduled, expired, or invalid

**Start Date**

Specifies the date and time when this rule will begin being used

**End Date**

Specifies the date and time when this rule will expire (no longer be used)

To change the name of the rule being copied, click on the Rule Name in the table and specify a new name for the rule.

Click **OK** when you are finished reviewing or changing the rule names. Click **Cancel** to abort the copy operation.

### View menu

This article describes the options available in the View menu window. The following options are available in this window:

## Status Bar

Use the **Status Bar** option to toggle the status bar on or off. The status bar is shown along the bottom of the Rule Browser window.

## Specify Columns

Use the **Specify Columns** option to adjust the type and order of the columns that display in your window.

In the window that opens, the following tasks can be accomplished:

### Add a new column

Select one or more entries in the **Available columns** list and click the **Add** button. The selected entries are added to the end of the **Columns displayed** list.

### Remove a column

Select one or more entries in the **Columns displayed** and click the **Remove** button. The selected entries are added to the end of the **Available columns** list.

### Reorder columns

Select one or more entries in the **Columns displayed** list. To move the entries towards the top of the list, click the Up arrow; to move the entries towards the bottom, click the Down arrow.

## Specify Date/Time Format

Use the **Specify Date/Time Format** option to adjust the format used when displaying dates and times. For more information, click the **Specify Date/Time Format window** link under Related reference.

## Refresh

Use the **Refresh** option to update the contents of the folder hierarchy and the rule table.

## Specify Date/Time Format window

Use the Specify Date/Time Format window to change the date and time format. Choose one of the following radio button options and then click **OK**:

**Use default format for this locale:** Use the **Use default format for this locale** option to adjust the date and time format to match the default setting of your current locale.

**Select a predefined format for the date and time:** Use the **Select a predefined format for the date and time** option to select one of several existing date and time formats.

**Specify a custom format for the date and time:** Use the **Specify a custom format for the date and time** option to determine your own format for your date and time display. Choose one of the date and time formats from the two menus.

## Find Rules window

Use the Find Rules window to specify search criteria to locate rules. The search combines your queries using a logical "AND" operation. For example, if you specify both a folder name and a rule name, the search finds rules that match both the folder name and rule name. It displays the results in a Search Results window.

**Note:** All of the fields in this window are optional.

## Main menu

The main menu has the following options:

- **File**

- Save As**

- Use the **Save as** option to open a Save Search window and store the current search criteria.

- Open** Use the **Open** option to open the Open Saved Search window and load a previously saved set of search criteria into the Find Rules window.

- Close** Use the **Close** option to close the Find Rules window.

- **View**

- Show Search**

- Use the **Show Search** option to display a text description of your search criteria on one screen.

- Specify Date/Time Format**

- Use the **Specify Date/Time Format** option to adjust the format used when displaying dates and times. For more information, see the **Specify Date/Time Format window** link in Related reference.

## **Tabs**

The following tabs are displayed in this window:

**Name** Use the **Name** tab to specify the basic search criteria. For more information, click the **Find Rules window: Name tab** link under Related reference.

**Date** Use the **Date** tab to specify the date-related search criteria. For more information, click the **Find Rules window: Date tab** link under Related reference.

**Classification**

Use the **Classification** tab to specify the search criteria related to a rule's classification. For more information, click the **Find Rules window: Classification tab** link under Related reference.

**Implementation**

Use the **Implementation** tab to specify the search criteria that is based on the manner in which the rule is implemented. For more information, click the **Find Rules window: Implementation tab** link under Related reference.

**Description**

Use the **Description** tab to specify the text-based search criteria related to a rule's description. For more information, click the **Find Rules window: Description tab** link under Related reference.

**Other** Use the **Other** tab to specify the search criteria based on precedence and user-defined data. For more information, click the **Find Rules window: Other tab** link under Related reference.

## **Find Rules window: Name tab**

Use the **Find Rules window: Name** tab to specify the basic search criteria. The following fields and options are available on this tab:

**Name:** Use the **Name** field to search for a specific rule name.

**Note:** This is case-sensitive.

**Drop-down search option list:** Use the one of the following options in the **Drop-down search option list** to narrow your search:

**equal** Use the **equal** selection to look for an *exact* match.

**starting with**

Use the **starting with** selection to find rules whose name *starts* with the specified value.

**ending with**

Use the **ending with** selection to find rules whose name *ends* with the specified value.

**containing**

Use the **containing** selection to find rules that *contain* the specified value.

**Location:** Use the **Location** field to specify the folder that you want to search. Click the **ellipses** button to the right of the field if you want to browse for the folder.

**Note:** The folder names are case-sensitive.

**Include subfolders:** Select **Include subfolders** to include the folder's subfolders in the search.

**Status:** Use the **Status** menu to specify a search criteria that is based on a rule's availability.

## Find Rules window: Date tab

Use the **Find Rules window: Date** tab to specify date-related search criteria. A rule always has a start date and a range of time in which it is in effect. The end date is optional and if it is not specified, the rule never expires. The following options are available on this tab:

**Find Rules for any date:** Use the **Find Rules for any date** option to remove the date from consideration in the search criteria.

**Find Rules that are:** Use the **Find Rules that are** option to search for rules in one of the following states:

### currently in effect

The **currently in effect** selection finds rules that are active at this point in time.

### scheduled

The **scheduled** selection finds rules that go into effect at a future date.

### expired

The **expired** selection finds previously active rules that are beyond the rule's end date.

**Find Rules:** Use the **Find Rules** option to specify the dates you want to search. Modify the following criteria to narrow your search:

1. Select one of the following rule states from the menu:

- in effect
- starting
- ending

2. Select **Query on date only** if you would like your search to ignore time-specific information.

3. Select one of the following three methods by which to search:

**on x** Select the **on x** option to find rules that are in the desired state (as chosen in Step 1) on the specified date. If the **Query on date only** check box is clear, then midnight is used for the time.

#### anytime between x and y

Select the **anytime between x and y** option to find rules that are in the desired state (as chosen in Step 1) *anytime* between the given dates. Specify a start date (represented by x) and an end date (represented by y). If the **Query on date only** check box is clear, then midnight of each day is used for the time.

#### anytime during the next x days

Select the **anytime during the next x days** option to find rules that are in the desired state (as chosen in Step 1) within this period of time and specified in days.

## Find Rules window: Classification tab

Use the **Find Rules window: Classification** tab to specify search criteria related to a rule's classification. The following options are available on this tab:

**Show all:** Use the **Show all** option if you do not want to include the classification information in the search criteria.

**Show rules that are not classified and do not perform classification:** Use the **Show rules that are not classified and do not perform classification** option to find rules that you do not need to classify.

**Show rules that perform classification:** Use the **Show rules that perform classification** option to find rules that return a classification such as "Gold", "Silver", or "Bronze".

**Show rules that are classified:** Use the **Show rules that are classified** option to find rules that are classified with the specified classification. The specified classification is case-sensitive. Enter a specific classification into the field that is provided.

## Find Rules window: Implementation tab

Use the **Find Rules window: Implementation** tab to specify search criteria that is based on the manner in which the rule is implemented. The following options are available on this tab:

**Java rule implementor:** Use the **Java rule implementor** option to search for rules that use the specified Java rule implementor. You can use one of the values in the list or type in your own.

**Firing location:** Use the **Firing location** option to search for the location in which the rule implementor is run. Choose one of the following values from the check box:

**Local** Use the **Local** option to search locally for the location in which the rule implementor is run.

**Remote**

Use the **Remote** option to search the server for the location in which the rule implementor is run.

**Anywhere**

Use the **Anywhere** option to search both locally and on the server for the location in which the rule implementor is run.

**Firing parameters:** Use the **Firing parameters** option to search for rules that alter the firing parameters passed from the trigger point to the Java rule implementor. For more information, see the **Add Firing Parameter window** link in Related reference.

The following selections are available for the firing parameter option:

- show rules that alter firing parameters
- show rules that do not alter firing parameters

### Find Rules window: Description tab

Use the **Find Rules window: Description** tab to specify text-based search criteria related to a rule's description. The following fields and options are available on this tab:

**Business intent:** Use the **Business intent** field to search for rules with a given business intent.

**Description:** Use the **Description** field to search for rules with a given description.

**Original requirement:** Use the **Original requirement** field to search for rules with a given original requirement.

**Drop-down search option list:** Use the following options in the **Drop-down search option list** to narrow your search in the **Business intent**, **Description**, and **Original requirements** fields:

**equal** Use the **equal** selection to look for an *exact* match.

**starting with**

Use the **starting with** selection to find rules whose name *starts* with the specified value.

**ending with**

Use the **ending with** selection to find rules whose name *ends* with the specified value.

**containing**

Use the **containing** selection to find rules that *contain* the specified value.

### Find Rules window: Other tab

Use the **Find Rules window: Other** tab to specify search criteria based on precedence and user-defined data. The following options are available on this tab:

**Precedence:** Use the **Precedence** option to search for rules with given precedence. The precedence is an integer value that specifies the relative priority of this rule when it is fired.

**Drop-down search option list:** Use the **Drop-down search option list** to narrow your search by using one of the following options:

- equal
- less than
- less than or equal
- greater than
- greater than or equal
- not equal

**User defined data:** Use the **User defined data** option to search for rules with given user defined data.

**Drop-down search option list:** Use one of the following options from the **Drop-down search option list**:

**equal** Use the **equal** selection to look for an *exact* match.

**starting with**

Use the **starting with** selection to find rules whose name *starts* with the specified value.

**ending with**

Use the **ending with** selection to find rules whose name *ends* with the specified value.

**containing**

Use the **containing** selection to find rules that *contain* the specified value.

## Search Results window

This window contains the results of a search from a Find Rules window. It is nearly identical to the Rule Browser in terms of form and function. Use the options in the main menu to perform many of the same administrative actions, with a few differences:

- Rules cannot be imported into the Search Results window.
- Rules cannot be pasted into the Search Results window.
- The **Open Containing Folder** menu item opens a new browser window showing the contents of the folder for the selected rule.

## Save Search window

Use the Save Search window to store the current search criteria for later retrieval. Follow the following steps to save your search:

1. Type in a name for your search or replace a previously saved search from the list.
2. Click **OK**.

The saved search criteria is loaded into the Find Rules window.

## Open Saved Search window

Use the Open Saved Search window to open a previously saved search. Follow these steps to open your saved search:

1. In the Select a Search window, highlight the name of the search you want to open.
2. Click **OK**.

The saved search criteria is loaded into the Find Rules window.

---

## Business rule beans: Resources for learning

Use the following links to find relevant supplemental information about business rule beans. The information resides on IBM and non-IBM Internet sites, whose sponsors control the technical accuracy of the information.

These links are provided for convenience. Often, the information is not specific to the IBM WebSphere Application Server product, but is useful all or in part for understanding the product. When possible, links are provided to technical papers and Redbooks that supplement the broad coverage of the release documentation with in-depth examinations of particular product areas.

View links to additional information about:

- “Business rule beans: Resources for learning”
- “Business rule beans: Resources for learning”
- “Business rule beans: Resources for learning”

## Planning, business scenarios, and IT architecture

- Delivering new business value to the enterprise on a J2EE and Web services base (Update)



This paper, in PDF form, describes the strategy behind the IBM extensions to J2EE and Web services functionality in the WebSphere Application Server Version 4.0 Enterprise Edition. It explains Enterprise Services, business rule beans, message beans and JMS listener, internationalization, shared work areas, bidirectional CORBA connectivity, the ActiveX bridge, and the C++ CORBA SDK.

- WebSphere Application Server V5.0 Architecture and Overview

This is an IBM developerWorks presentation that provides an overview of the functionality available with WebSphere Application Server V5.0.

### **Programming instructions and examples**

- Message-Driven Beans and Encapsulated Business Rules

This article describes how to use business rules with Message-driven Beans.

- WebSphere Application Server Enterprise Edition 4.0: A Programmer's Guide

Chapter three of this programmer's guide provides information about implementation, modification, and deployment of business rules.

- WebSphere Application Server Enterprise Edition Technology Sample

This technology sample enables developers to gain experience with the Business Rule Beans technology.

### **Administration**

- IBM WebSphere Administration
- Listing of all IBM WebSphere Application Server Redbooks

This is a listing of the Redbook publications about the WebSphere software platform.

- WebSphere Application Server Version 4.0 Enterprise Edition -- Presentations and Labs



---

## Chapter 18. Using asynchronous beans

The asynchronous beans feature adds a new set of APIs that enable Java 2 Platform Enterprise Edition J2EE applications to run asynchronously inside an Integration Server. This topic provides a brief overview of the tasks involved in using asynchronous beans. For a more detailed description of the asynchronous beans model, review the conceptual topic Asynchronous beans.

1. Configure work managers.
2. Assemble applications that use work managers.
3. Develop work objects to run code in parallel.
4. Develop event listeners.
5. Develop asynchronous scopes.

---

### Asynchronous beans

An asynchronous bean is a Java object or enterprise bean that can be executed asynchronously by a Java 2 Platform Enterprise Edition (J2EE) application, using the J2EE context of the asynchronous bean creator.

Asynchronous beans can improve performance by enabling a J2EE program to decompose operations into parallel tasks. Asynchronous beans support the construction of stateful, active J2EE applications. These applications address a segment of the application space that J2EE has not previously addressed (that is, advanced applications that require application threading, active agents within a server application, or distributed monitoring capabilities).

Asynchronous beans can run using the J2EE security context of the creator J2EE component. These beans also can run with copies of other J2EE contexts. For example:

- Internationalization context
- Application profiles
- Work areas

### Asynchronous bean interfaces

Three types of asynchronous bean exist:

#### Work object

A work object implements the `com.ibm.websphere.asynchbeans.Work` interface. A work object runs parallel to its caller using the `WorkManager.startWork()` method. Applications implement work objects to run code blocks asynchronously. For more information on the `Work` interface, see the Related reference section at the end of this article.

#### Alarm listener

An alarm listener is an object that implements the `com.ibm.websphere.asynchbeans.AlarmListener` interface. Alarm listeners are called when a high-speed transient alarm expires. For more information on the `AlarmListener` interface, see the Related reference section at the end of this article.

#### Event listener

An event listener can implement any interface. An event listener is a lightweight, asynchronous notification mechanism for asynchronous events within a single Java virtual machine (JVM). An event listener typically enables J2EE components within a single application to notify each other about various asynchronous events.

### Supporting interfaces

#### Work manager

A work manager is a thread pool that administrators create for J2EE applications. The administrator specifies the properties of the thread pool and a policy that determines which J2EE contexts the asynchronous bean inherits.

## Event source

An event source implements the `com.ibm.websphere.asynchbeans.EventSource` interface. An event source is a system-provided object that supports a generic, type-safe asynchronous notification server within a single JVM. The event source enables event listener objects, which implement any interface to be registered. For more information on the `EventSource` interface, see the Related reference section at the end of this article.

## Event source events

Every event source can generate events of its own. Event sources also can generate their own events such as listener count changed. An application can register an event listener object that implements the class `com.ibm.websphere.asynchbeans.EventSourceEvents`. This action enables the application to catch events such as listeners being added or removed, or a listener throwing an unexpected exception. For more information on the `EventSourceEvents` class, see the Related reference section at the end of this article.

Additional interfaces, including alarms and subsystem monitors, are introduced in the topic *Developing Asynchronous scopes*, which discusses some of the advanced applications of asynchronous beans.

## Transactions

Every asynchronous bean method is called using its own transaction, much like container-managed transactions in typical enterprise beans. It is very similar to the situation when an Enterprise Java Beans (EJB) method is called with `TX_NOT_SUPPORTED`. The run-time environment starts a local transaction before invoking the method. The asynchronous bean method is free to start its own global transaction if this transaction is possible for the calling J2EE component. For example, if an enterprise bean creates the component, the method that creates the asynchronous bean must be `TX_BEAN_MANAGED`.

When you call an entity bean from within an asynchronous bean, for example, you must have a global transactional context available on the current thread. Because asynchronous bean objects start local transactional contexts, you can encapsulate all entity bean logic in a session bean that has a method marked as `TX_REQUIRES` or equivalent. This process establishes a global transactional context from which you can access one or more entity bean methods.

If the asynchronous bean method throws an exception, any local transactions are rolled back. If the method returns normally, any incomplete local transactions are completed according to the unresolved action policy configured for the bean. EJB methods can configure this policy using their deployment descriptor. If the asynchronous bean method starts its own global transaction and does not commit this global transaction, the transaction is rolled back when the method returns.

## Access to J2EE component meta data

If an asynchronous bean is a J2EE component, such as a session bean, its own meta data is active when a method is called. If an asynchronous bean is a simple Java object, the J2EE component metadata of the creating component is available to the bean. Like its creator, the asynchronous bean can look up the `java:comp` namespace. This look up enables the bean to access connection factories and enterprise beans, just as it would if it were any other J2EE component. The environment properties of the creating component also are available to the asynchronous bean.

The `java:comp` namespace is identical to the one available for the creating component; the same restrictions apply. For example, if the enterprise bean or servlet has an EJB reference of `java:comp/env/ejb/MyEJB`, this EJB reference is available to the asynchronous bean. In addition, all of the connection factories use the same resource-sharing scope as the creating component.

## Connection management

An asynchronous bean method can use the connections that its creating J2EE component obtained using `java:comp` resource references. (For more information on resource references, see *References*). However,

the bean method must access those connections using a get, use, close pattern. There is no connection caching between method calls on an asynchronous bean. The connection factories or datasources can be cached, but the connections must be retrieved on every method call, used, and then closed. While the asynchronous bean method can look up connection factories using a global Java Naming and Directory Interface (JNDI) name, this is not recommended for the following reasons:

- The JNDI name is hard coded in the application (for example, as a property or string literal).
- The connection factories are not shared because there is no way to specify a sharing scope.

For code examples that demonstrate both the correct and the incorrect ways to access connections from asynchronous bean methods, see the topic [Example: Asynchronous bean connection management](#).

## Deferred start of Asynchronous Beans

Asynchronous Beans supports deferred execution by allowing serialization of J2EE service context information. The `WorkWithExecutionContext` `createWorkWithExecutionContext(Work r)` method on the `WorkManager` interface will create a snapshot of the J2EE service contexts enabled on the `WorkManager`. The resulting `WorkWithExecutionContext` object can then be serialized and stored in a database or file. This is useful when it is necessary to store J2EE service contexts such as the current security identity or `Locale` and later inflate them and execute some work within this context. The `WorkWithExecutionContext` object can be executed using the `startWork()` and `doWork()` methods on the `WorkManager` interface.

All `WorkWithExecutionContext` objects must be deserialized by the same application that serialized it. All EJBs and classes must be present in order for Java to successfully inflate the objects contained within.

## Deferred Execution and Security

The asynchronous beans security service context might require Common Secure Interoperability Version 2 (CSIv2) identity assertion to be enabled. Identity assertion is required when a `WorkWithExecutionContext` object is deserialized and executed to Java Authentication and Authorization Service (JAAS) subject identity credential assignment. Review the following topics in the Information Center to better understand if you need to enable identity assertion, when using a `WorkWithExecutionContext` object:

- [Configuring Common Secure Interoperability Version 2 and Security Authentication Service authentication protocol](#)
- [Identity Assertion](#)

There are also issues with interoperating with `WorkWithExecutionContext` objects from different versions of the product. See [Interoperating with asynchronous beans](#).

## Example: Asynchronous bean connection management

An asynchronous bean method can use the connections that its creating Java 2 Platform Enterprise Edition (J2EE) component obtained using `java:comp` resource references. (For more information on resource references, see the topic [References](#).) The following is an example of an asynchronous bean that uses connections correctly:

```
class GoodAsynchBean
{
 DataSource ds;
 public GoodAsynchBean()
 throws NamingException
 {
 // ok to cache a connection factory or datasource
 // as class instance data.
 InitialContext ic = new InitialContext();
 // it is assumed that the created J2EE component has this
 // resource reference defined in its deployment descriptor.
 ds = (DataSource)ic.lookup("java:comp/env/jdbc/myDataSource");
 }
}
```

```

// When the asynchronous bean method is called, get a connection,
// use it, then close it.
void anEventListener()
{
 Connection c = null;
 try
 {
 c = ds.getConnection();
 // use the connection now...
 }
 finally
 {
 if(c != null) c.close();
 }
}
}

```

The following example of an asynchronous bean that uses connections illegally:

```

class BadAsynchBean
{
 DataSource ds;
 // Do not do this. You cannot cache connections across asynch method calls.
 Connection c;

 public BadAsynchBean()
 throws NamingException
 {
 // ok to cache a connection factory or datasource as
 // class instance data.
 InitialContext ic = new InitialContext();
 ds = (DataSource)ic.lookup("java:comp/env/jdbc/myDataSource");
 // here, you broke the rules...
 c = ds.getConnection();
 }
 // Now when the asynch method is called, illegally use the cached connection
 // and you likely see J2C related exceptions at run time.
 // close it.
 void someAsynchMethod()
 {
 // use the connection now...
 }
}

```

---

## Configuring work managers

If you are not familiar with work managers, review the conceptual topic [Work managers](#).

A work manager acts as a thread pool for application components that use asynchronous beans. Use the administrative console to configure work managers. You can define multiple work managers for each cell. Each work manager is bound to a unique place in Java Naming and Directory Interface (JNDI).

**Note:** The work manager service is only supported from within the Enterprise Java Beans (EJB) Container or Web Container. Looking up and using a configured work manager from a J2EE application client container is not supported.

1. Start the administrative console.
2. Select **Resources > Work Managers**.
3. Click **New**.
4. Specify the following required properties:
  - Name** The display name for the work manager.

**JNDI Name**

The JNDI name for the work manager. This name is used by asynchronous beans that need to look up the work manager. Each work manager must have a unique JNDI name within the cell.

**Number of Alarm Threads**

The maximum number of threads to use for processing alarms. A single thread is used to monitor pending alarms and dispatch them. An additional pool of threads is used for dispatching the threads. All alarm managers on the asynchronous beans associated with this work manager share this set of threads. A single alarm thread pool exists for each work manager, and all of the asynchronous beans associated with the work manager share this pool of threads.

**Minimum Number Of Threads**

The initial number of threads to be created in the thread pool.

**Maximum Number Of Threads**

The maximum number of threads to be created in the thread pool. The maximum number of threads can be exceeded temporarily if the **Growable** check box is selected. These additional threads are discarded when the work on the thread completes.

**Thread Priority**

The order of the priority for threads available in the thread pool.

5. [Optional] Specify a **Description** and a **Category** for the work manager.
6. [Optional] Select the **Service Names** (J2EE contexts) on which you want this work manager to be made available. Any asynchronous beans that use this work manager then inherit the selected J2EE contexts from the component that creates the bean. The list of selected services also is known as the "sticky" context policy for the work manager. Selecting more services than are actually required might impede performance.
7. Save your configuration.

The work manager is now configured and ready for access by application components that need to manage asynchronous code execution.

## Work managers

A work manager is a thread pool created for J2EE applications that use asynchronous beans.

Using the administrative console, an administrator can configure any number of work managers. The administrator specifies the properties of the work manager, including the inheritance context policy for any asynchronous beans that use the work manager. The administrator binds each work manager to a unique place in Java Naming and Directory Interface (JNDI).

When writing a Web or EJB component that uses asynchronous beans, the developer should include a resource reference in each component that needs access to a work manager. (For more information on resource references, see the topic [References](#).) The component looks up a work manager using a logical name in the component's `java:comp` namespace, just as it looks up a data source, enterprise bean, or connection factory.

The deployer binds physical work managers to logical work managers when the application is deployed.

For example, if a developer needs three thread pools to partition work between bronze, silver, and gold levels, the developer writes the component to pick a logical pool based on an attribute in the client application profile. The deployer has the flexibility to decide how to map this request for three thread pools. The deployer might decide to use a single thread pool on a small machine. In this case, the deployer binds all three resource references to the same work manager instance (that is, the same JNDI name). A larger machine might support for three thread pools, so the deployer binds each resource reference to a different work manager. Work managers can be shared between multiple J2EE applications installed on the same server.



An application developer can use as many logical work managers as necessary; the deployer chooses whether to map one physical work manager or several to the logical work manager defined in the application.

All J2EE components that need to share asynchronous scope objects must use the same work manager. These scope objects have an affinity with a single work manager so an application that uses `AsynchScopes` should verify that all of the components using scope objects use the same work manager.

When multiple work managers are defined, the underlying thread pools are created in a JVM only if an application within that Java virtual machine (JVM) looks up the work manager. For example, there might be ten thread pools (work managers) defined, but none are actually created until an application looks these pools up.

### How to look up a work manager

An application can look up a work manager as follows. Here, the component contains a resource reference named `wm/myWorkManager`, which was bound to a physical work manager when the component was deployed:

```
InitialContext ic = new InitialContext();
WorkManager wm = (WorkManager)ic.lookup("java:comp/env/wm/myWorkManager");
```

### Inheritance J2EE contexts

Asynchronous beans can inherit the following J2EE contexts.

#### Internationalization context

#### Work area

#### Application profile

#### Security

The asynchronous bean can be run as anonymous or as the client authenticated on the thread that created it. This is useful because the asynchronous bean can do only what the caller can do.

This is more useful than a `RUN_AS` mechanism, for example, which prevents this kind of behavior.

#### Component meta data

Component meta data is relevant only when the asynchronous bean is a simple Java object. If the bean is a J2EE component, such as an enterprise bean, the component's meta data is active.

The contexts that can be inherited depends on the work manager used by the application that creates the asynchronous bean. Using the administrative console, the administrator defines the sticky context policy of a work manager by selecting the services on which the work manager is to be made available.

## Work manager collection

Use this page to view the configuration properties of work managers.

A work manager contains a pool of threads bound into JNDI.

To view this administrative console page, click **Resources > Work Managers** .

### Name

The name by which the work manager is known for administrative purposes.

**Data type** String

### JNDI Name

The JNDI name used to look up the work manager in the namespace.

**Data type** String

## Service Names

A list of service names on which this work manager is made available.

The context information of each selected service is propagated to each work or alarm that is created using this work manager. Selecting services that are not needed can negatively impact performance.

**Data type** Integer

## Number of Alarm Threads

The number of threads used to execute concurrent alarms.

**Data type** Integer

## Minimum Number of Threads

The minimum number of threads available in this work manager for running works.

**Data type** Integer

## Maximum Number of Threads

The maximum number of threads available in this work manager for running works.

**Data type** Integer

## Work manager settings

Use this page to modify work manager settings.

A work manager contains a pool of threads bound into JNDI.

To view this administrative console page, click **Resources > Work Managers > *workmanager\_name***.

### **Name:**

The name by which the work manager is known for administrative purposes.

**Data type** String

### **JNDI Name:**

The JNDI name used to look up the work manager in the namespace.

**Data type** String

### **Description:**

A description of this work manager for administrative purposes.

**Data type** String

### **Category:**

A string that can be used to classify or group this work manager.

**Data type** String

***Number of Alarm Threads:***

The number of threads used to execute concurrent alarms.

**Data type** Integer

***Minimum Number of Threads:***

The minimum number of threads available in this work manager for running works.

**Data type** Integer

***Maximum Number of Threads:***

The maximum number of threads available in this work manager for running works.

**Data type** Integer

***Thread Priority:***

The priority of the threads available in this work manager.

**Data type** Integer

***Growable:***

Specifies whether the number of threads in this work manager can be increased.

***Service Names:***

A list of service names on which this work manager is made available.

The context information of each selected service is propagated to each work or alarm that is created using this work manager. Selecting services that are not needed can negatively impact performance.

**Data type** Integer  
**Range** Application Profiling Service  
Work Area  
Security  
Internationalization

## **Work manager service settings**

Use this page to enable or disable the work manager service, which manages work manager resources used by the server.

To view this administrative console page, click **Servers > Application Servers > *server\_name* > Work Manager Service** .

### **Startup**

Specifies whether the server will attempt to start the work manager service.

**Default  
Range**

**Selected  
Selected**

When the application server starts, it attempts to start the work manager service automatically.

**Cleared**

The server does not try to start the work manager service. If work manager resources are to be used on this server, the system administrator must start the work manager service manually or select this property then restart the server.

---

## Assembling applications that use work managers

Your administrator needs to configure at least one work manager using the administrative console.

If your application references one or more logical work managers, the logical work managers must be bound to one or more physical work managers using the Application Assembly Tool (AAT) Assembly Toolkit.

1. Declare a resource reference for each work manager (required action by the application developer). This forms an EAR file. (For more information on resource references, see the topic [References](#).)
2. Using the Application Assembly Tool (AAT) Assembly Toolkit, bind each resource reference to a physical work manager.
3. Add a resource reference with the type 'com.ibm.websphere.asynchbeans.WorkManager' to the application's descriptor. The application then can look up this work manager using its resource reference name in java:comp. The assembly tool or WebSphere Studio Application Developer Integration Edition then can specify which resource references are bound to a physical work manager.

**Note:** The previous steps is the same process used for data sources.

---

## Developing work objects to run code in parallel

Your administrator must have configured at least one work manager using the administrative console.

To run code in parallel, or in a different J2EE context, wrap the code in a work object.

1. Create a work object.

A work object implements the `com.ibm.websphere.asynchbeans.Work` interface. For example:

```
class SampleWork implements Work
```

2. Determine the number of work managers needed by this application component.
3. Look up the work manager or managers using the work manager resource reference (or logical name) in the `java:comp` namespace. (For more information on resource references, see the topic [References](#).)

```
InitialContext ic = new InitialContext();
WorkManager wm = (WorkManager)ic.lookup("java:comp/env/wm/myWorkManager");
```

The resource reference for the work manager (in this case, `wm/myWorkManager`) must be declared as a resource reference in the application deployment descriptor.

4. Call the `WorkManager.startWork()` method using the work object as a parameter. For example:

```
Work w = new MyWork(...);
WorkItem wi = wm.startWork(w);
```

The `startWork()` method can take a `startTimeout` parameter. This specifies a hard time limit for the Work object to be started. The `startWork()` method returns a work item object. This object is a handle that provides a link from the component to the now running work object.

5. [Optional] If your application component needs to wait for one or more of its running work objects to complete, call the `WorkManager.join()` method. For example:

```
WorkItem wiA = wm.start(workA);
WorkItem wiB = wm.start(workB);
ArrayList l = new ArrayList();
l.add(wiA);
l.add(wiB);
if(wm.join(l, wm.JOIN_AND, 5000)) // block for up to 5 seconds
{

 // both wiA and wiB finished
}
else
{

 // timeout

 // we can check wiA.getStatus or wiB.getStatus to see which, if any, finished.
}
```

This method takes an array list of work items which your component wants to wait on and a flag that indicates whether the component will wait for one or all of the work objects to complete. You also can specify a timeout value.

6. Use the `release()` method to set a variable in a synchronized block. For example:

```
public synchronized void release()
{
 released = true;
}
```

The `Work.run()` method can periodically examine this variable to check whether the loop exits or not.

## Work objects

A work object is a type of asynchronous bean used by application components to run code in parallel or in a different J2EE context.

A work object implements the `com.ibm.websphere.asynchBeans.Work` interface. A work object is essentially a `java.lang Runnable` object that is serializable and provides additional methods. For details, see the Interface `Work` in the Javadoc.

A component wanting to run work in parallel, or in a different J2EE context, locates a work manager in JNDI, then calls the `WorkManager.startWork()` method using the work object as a parameter.

The `startWork()` method returns a work item object. This object is a handle that provides a link from the component to the now running work object. The work item object is typically used when the component needs to wait for one or more of its running work objects to complete. The `WorkManager.join()` method takes an array list of work items that the component wants to wait on, and a flag indicating whether the component will wait for all or one of the work objects to complete. A timeout can be specified, which prevents the component from waiting indefinitely.

Why not have the application create Java 2 SDK threads? The threads created by the Java 2 SDK are not managed threads and hence know nothing about the J2EE environment and are unusable inside an application server. In addition, these threads have no J2EE context (for example, a `java:comp`) and are not authenticated when they fire. Work object threads are fully supported by the application server and have the same properties as any other asynchronous bean.

## Example: Work object

The following is an example of a work object that dynamically subscribes to a topic:

```
class SampleWork implements Work
{
 boolean released;
 Topic targetTopic;
 EventSource es;
 TopicConnectionFactory tcf;
 public SampleWork(TopicConnectionFactory tcf, EventSource es, Topic targetTopic)
 {
 released = false;
 this.targetTopic = targetTopic;
 this.es = es;
 this.tcf = tcf;
 }
 synchronized boolean getReleased()
 {
 return released;
 }
 public void run()
 {
 try {
 // setup our JMS stuff.
 TopicConnection tc = tcf.createConnection();
 TopicSession sess = tc.createSession(false, Session.AUTOACK);
 tc.start();

 MessageListener proxy = es.getEventTrigger(MessageListener.class, false);
 while(!getReleased())
 {
 // block for up to 5 seconds.
 Message msg = sess.receiveMessage(5000);
 if(msg != null)
 {
 // fire an event when we get a message
 proxy.onMessage(msg);
 }
 }
 tc.close();
 } catch (JMSEException ex)
 {
 // handle the exception here
 throw ex;
 }
 finally
 {
 if (tc != null)
 {
 try{
 tc.close();
 } catch (JMSEException ex1)
 { // handle exception }
 }
 }
 }
 // called when we want to stop the Work object.
 public synchronized void release()
 {
 released = true;
 }
}
```

As a result, any component that has access to the event source can add an event on demand, which allows components to subscribe to a topic in a more scalable way than by simply giving each client subscriber its own thread. The previous example is fully explored in the WebSphere Trader Sample. See your Samples Gallery for details.

---

## Developing event listeners

Application components that listen for events can use the `EventSource.addListener()` method to register an event listener object (a type of asynchronous bean) with the event source to which the events will be published. An event source also can fire events in a type-safe manner using any interface.

Notifications between components within a single EAR file are handled by a special event source. See the topic, [Using the application notification service](#).

1. Create an event listener object, which can be any type. For example, see the following interface code:

```
interface SampleEventGroup
{
 void finished(String message);
}

class myListener implements SampleEventGroup
{
 public void finished(String message)
 {
 // This will be called when we 'finish'.
 }
}
```

2. Register the event listener object with the event source. For example, see the following code:

```
InitialContext ic = ...;
EventSource es = (EventSource)ic.lookup("java:comp/websphere/ApplicationNotificationService");
myListener l = new myListener();
es.addListener(l);
```

This enables the `myListener.finished()` method to be called whenever the event is fired. The following code example shows how this event might be fired:

```
InitialContext ic = ...;
EventSource es = (EventSource)ic.lookup("java:comp/websphere/ApplicationNotificationService");
myListener proxy = es.getEventTrigger(myListener.class);
// fire the 'event' by calling the method
// representing the event on the proxy
proxy.finished("done");
```

## Using the application notification service

During the application lifetime, individual J2EE components (servlets or enterprise beans) within a single EAR file might need to signal each other. There is an event source in the `java:comp` namespace that is bound into all components within an EAR file. The JNDI name for this event source is:

```
java:comp/websphere/ApplicationNotificationService
```

Components within the same application can fire asynchronous events and register event listeners using this application notification service. Startup beans can be used to register these event listeners at application startup or they can be registered dynamically at run time.



To have your enterprise bean or servlet use the application notification service, write code similar to the following example:

```
InitialContext ic = new InitialContext();
EventSource appES = (EventSource)
 ic.lookup("java:comp/websphere/ApplicationNotificationService");
// now, the application can add a listener using the EventSource.addListener method.
// MyEventType is an interface.
MyEventType myListener = ...;
AppES.addListener(myListener);

// later another component can fire events as follows
InitialContext ic = new InitialContext();
EventSource appES = (EventSource)
ic.lookup("java:comp/websphere/ApplicationNotificationService");

// This highlights a constant string on the EventSource interface which
// specifies the 'java:comp/websphere/ApplicationNotificationService' string.
ic.lookup(appES.APPLICATION_NOTIFICATION_EVENT_SOURCE)
// now, the application can add a listener using the EventSource.addListener method.
MyEventType proxy = appES.getEventTrigger(MyEventType.class, false);
proxy.someEvent(someArguments);
```

## Example: Event listener

The following code example demonstrates how to fire a listenerCountChanged event:

```
// imagine this snippet inside an EJB or servlet method.
// Make an inner class implementing the required event interfaces.
EventSourceEvents listener = new Object() implements EventSourceEvents.class
{
 void listenerCountChanged(EventSource es, int old, int newCount)
 {
 try
 {
 InitialContext ic = new InitialContext();
 // here, the asynch bean can access an environment variable of
 // the component which created it.
 int i = (Integer)ic.lookup("java:comp/env/countValue").intValue();
 if(newCount == i)
 {
 // do something interesting
 }
 // call this event when the following code executes:
 }
 catch(NamingException e)
 {
 }
 }
 void listenerExceptionThrown(EventSource es, Object listener,
 String methodName, Throwable exception)
 {
 }
 void unexpectedException(EventSource es, Object runnable, Throwable exception)
 {
 }
}
// register it.
es.addListener(listener);

...

// now fire an event which the previous listener receives.
EventSourceEvents proxy = (EventSourceEvents)
 es.getEventTrigger(EventSourceEvents.class, false);
```

```
proxy.listenerCountChanged(es, 0, 1);

// now, fire another event, you can call any of the methods.
proxy.listenerCountChanged(es, 4, 5);
```

In this example, you get a proxy for the interface that we want to fire a method on. Then, call the method corresponding to the event on the proxy. This action causes the same method with the same parameters to be called on any event listeners that implement the `EventSourceEvents` interface and that were previously registered with the EventSource "es". The same proxy can be used to send multiple events simultaneously.

The boolean parameter on the `getEventTrigger()` method is named "sameTransaction". When the `sameTransaction` parameter is `false`, a new transaction is started for each event listener invoked and these event listeners can be called in parallel to the caller. However, the `event()` method always is blocked until all of the event listeners are notified. If the `sameTransaction` parameter is `true`, the current transaction, if any, on the thread is used for all of the event listeners; that is, the event listeners share the transaction of the method that fired the event. For that reason, all event listeners must run serially in an undetermined order. That is, the order in which the listeners are called is undefined and the order in which listeners are registered does not act as a guide for the order used at run time. The method on the proxy does not return until all of the event listeners are called; that is, it is a synchronous operation.

The parameters that are passed by reference and listeners do not interfere with these references unless the method that fired the event has purposefully designed such interaction. For example, event listeners can be used as collaborators and add data to a map, which was a parameter. Each event listener runs on its own transaction, independent of any transaction that is active on the thread. Extreme care must be taken when the `sameTransaction` parameter is `false` because the parameters can be accessed by multiple threads.

---

## Developing asynchronous scopes

Asynchronous scopes are units of scoping that comprise a set of alarms, subsystem monitors, and child asynchronous scopes. Using asynchronous scopes can involve some or all of the following steps:

1. Create asynchronous scopes Create a parent asynchronous scope object by calling the `AsynchScopeManager.createAsynchScope()` method using a unique name as the parameter. You can store properties in an asynchronous scope object. This storage provides Java 2 Enterprise Edition (J2EE) applications with a way to store a non-serializable state that otherwise cannot be stored in a session bean. You also can create child asynchronous scopes, which is useful for scoping data beneath the parent.
2. Listen for alarm notifications
  - a. Create a listener object by implementing the `AlarmListener` interface. For more information, see the `AlarmListener` interface in the Javadoc.
  - b. Supply this object to the `AlarmManager.create()` method, as the target for the alarm. The `create()` method takes the following parameters:

### Target for the alarm

The target on which the `fired()` method is called when the alarm is fired.

### Context

The context object for the alarm. This object is useful for supplying alarm-specific data to the listener and supports a single listener for multiple alarms.

### Interval

The number of milliseconds before the alarm fires.

After the specified interval, the alarm fires and the `fired()` method of the listener is called with the firing alarm as a parameter. The alarm object is returned. By calling methods on this object, you can cancel or reschedule the alarm.

3. Monitor remote systems.

- a. Implement a mechanism for detecting messages sent from the remote system. For example, publish and subscribe messaging.
- b. Create a subsystem manager object by calling the `SubsystemMonitorManager.create()` method with the following parameters:
  - Name** Each subsystem monitor must have a unique name.
  - Heartbeat interval**  
The expected interval, in milliseconds, between heartbeats.
  - Missed heart beats until stale or suspect**  
The number of heartbeats that can be missed before the subsystem is marked as stale.
  - Missed heart beats until dead**  
The number of heartbeats that can be missed before the system is marked as dead.
- c. Create an object that implements the `SubsystemMonitorEvents` interface. For more information, see the `SubsystemMonitorEvents` in the Javadoc.
- d. Add an instance of this object to the subsystem monitor using the `SubsystemMonitor.addListener()` method.
- e. Whenever a heartbeat message arrives from the remote system, call the `SubsystemMonitor.ping()` method.

The subsystem monitor configures alarms to track the heartbeat status of the remote system. When the `ping()` method is called, the alarms are reset. If an alarm fires, the `ping()` method is not called; that is, the application did not receive a heartbeat from the monitored subsystem.

Asynchronous scopes are useful in stateful server applications. An application can have a startup bean that creates an asynchronous scope on a named work manager. The application also might create subsystem monitors to monitor the health of any remote systems on which the application is dependent.

When a client attaches to the server, the application creates a child asynchronous scope that is owned by the application asynchronous scope for the client and named using the client ID. A subsystem monitor for monitoring the client might be created on the client asynchronous scope. If the client times out, a callback can clean up the client state on the server. Callbacks can be attached to the application subsystem monitors, on behalf of the client. When a remote system becomes unavailable, the client code in the server is notified and an event is sent to the client to warn that a critical remote system has failed. For example, the failure might be a data feed in an electronic trading application.

## Asynchronous scopes

An asynchronous scope (`AsynchScope` object) is a unit of scoping provided for use with asynchronous beans.

Asynchronous scopes are collections of alarms, subsystem monitors, and child asynchronous scopes that enable a relationship to form. Each asynchronous scope uses a single work manager.

Each `AsynchScope` object owns and controls the life cycle of the following objects:

### Child asynchronous scopes

Each `AsynchScope` object extends the `AsynchScopeManager` interface, which is a factory for `AsynchScope` objects. (For more information on the `AsynchScopeManager` interface, see the Javadoc). Any asynchronous scope can therefore create named asynchronous scopes (children). Child asynchronous scopes can be useful for scoping data underneath the parent. All of the child asynchronous scopes must be uniquely named. These children are destroyed if the parent asynchronous scope is destroyed.

### Alarms

Each asynchronous scope has an associated alarm manager. All of the alarms created by the alarm manager are automatically cancelled if the associated asynchronous scope is destroyed.

### Subsystem monitors

Each asynchronous scope has a subsystem monitor manager, which manages a set of subsystem

monitors associated with the asynchronous scope. When the asynchronous scope is destroyed, all of the associated subsystem monitors also are destroyed.

To summarize, asynchronous scopes can be organized into an acyclic tree. The life cycle of each asynchronous scope is directly coupled to that of its parent asynchronous scope. Each asynchronous scope is associated with a set of alarms and subsystem monitors, and an optional set of child asynchronous scopes. These objects are cancelled and destroyed when the asynchronous scope is destroyed.

### **Asynchronous scope state**

Each asynchronous scope has an associated map, in which applications can store their state in the form of name and value pairs.

### **Asynchronous scope events**

Each asynchronous scope is also an event source. Applications can therefore register event listeners against the asynchronous scope. The event listeners can receive notification if, for example, the `AsynchScope` object is about to be destroyed.

Applications also can use this event source to fire events only to listeners of this asynchronous scope. For example, an `AsynchScope` object created for a client session might be used to fire asynchronous events to parties interested in that client.

## **Alarms**

An alarm executes Java 2 Enterprise Edition (J2EE) context-aware code at a given time interval. Alarm objects are fine-grained, non-persistent, transient, and can fire at millisecond intervals.

Alarms are executed using a thread pool associated with the work manager that owns the associated asynchronous scope.

The `AlarmManager.createAlarm()` method takes an application-written object that implements the `AlarmListener` interface. (For more information on the `AlarmListener` interface, see the Javadoc) The `fire` method is called when the alarm expires. The `createAlarm()` method returns a non-serializable handle, which can be used to cancel or reset the alarm. All of the pending alarms are cancelled when its associated `AsynchScope` object is destroyed.

The Java 2 Software Development Kit (SDK) already has a timer mechanism, so why create a new one? The Java 2 SDK is a Java 2 Platform Standard Edition (J2SE) feature that knows nothing about the J2EE environment. Timers fired by the J2SE feature do not run on a managed thread and are therefore unusable inside an application server. These timers also lack a J2EE context (that is, a `java:comp` value) and are not authenticated when they fire. The asynchronous scope alarms are fully supported by the product and have the same properties as any other asynchronous bean.

### **Alarm performance**

The alarm subsystem is designed to handle a large number of alarms. However, do not have alarms undertake heavy processing when they are firing because this activity slows the processing of later alarms. If an alarm needs to process a heavy load, design a work object that is activated by a work manager. This procedure moves the heavy processing to a different thread and enables the alarm threads to process alarms unhampered. All of the alarms owned by asynchronous scopes that are owned by a single work manager, share a common thread pool. The properties of this thread pool can be tuned at the work manager level using the administrative console.

## Subsystem monitors

A subsystem monitor is an object that monitors the health of a remote system. It uses an event source to inform all registered listeners of the health of the system.

Advanced Java 2 Platform Enterprise Edition (J2EE) applications often rely on remote, non-managed, non-J2EE systems. These remote systems can periodically send clients a message to indicate that they are working. A subsystem monitor is essentially a set of alarms that track indicators messages or heartbeats from a remote system.

An application creates a subsystem monitor by calling the `SubsystemMonitorManager.create()` method with the following parameters:

**Name** Each subsystem monitor must be uniquely named.

**Heart beat interval**

The time period, in milliseconds, between arriving heartbeat messages.

**Missed heart beats until stale or suspect**

The number of heartbeats that can be missed before the subsystem is marked as stale. This designation indicates that the subsystem might be having problems.

**Missed heart beats until dead**

The number of heartbeats that can be missed before the system is considered down. The system then is marked as dead.

The subsystem monitor configures alarms to track the heartbeat status. Whenever the `ping()` method is called, the alarms are reset. If an alarm fires, the `ping()` method has not been called; that is, the application did not receive a heartbeat from the monitored subsystem. When the number of **Missed heart beats until stale** value has elapsed without a ping, a stale event is fired. Later, if the number of **Missed heart beats until dead** value elapses without a ping, a dead event is fired. If a ping is received after a stale or dead notification, a fresh event is sent, which indicates that the subsystem is alive again.

Make the **Missed heart beats until dead** value greater or equal to the **Missed heart beats until stale** value. If **Missed heart beats until stale** value equals the **Missed heart beats until dead** value, then a stale event is not published. Only a dead event is published.

Applications wanting to be informed of these events can register a listener that implements the `SubsystemMonitorEvents` interface. For more information on the `SubsystemMonitorEvents` interface, see the Javadoc.

Heart beat messages can be transmitted using a variety of mechanisms. The application must call the `SubsystemMonitor ping()` method whenever a heartbeat message arrives from a remote system, but the method used to detect these messages is up to the application. For example, you might use a Java Message Service (JMS) publish or subscribe implementation or even a third-party Java messaging product that does not implement JMS.

## Asynchronous scopes: Dynamic message bean scenario

Java 2 Platform Enterprise Edition (J2EE) now supports message-driven beans, but the beans are static. All of the message sources must be known in advance and bound at deployment time. This action is not always viable, especially in fluid messaging environments such as those found in brokerages. Some environments have publish-subscribe topic spaces that are continually changing and clients need servers that can subscribe on demand to an arbitrary topic.

An asynchronous bean application can create a work object that performs a blocking receive on a Java Messaging Service (JMS) topic and then publishes the message as an event on an application-defined event source. Clients requiring a subscription to that message can add an event listener to the event source. The event source can inform the work object when there are no listeners. Then, the event source can shut down and make the JMS and thread resources available. The work object registers a listener with its own event source. When the count is one again, the work object knows that it is the only listener

and it is time to shut down the work object. The WebSphere Trader Sample (see your installed Samples Gallery) uses this pattern to dynamically subscribe to JMS topics at run time to gather stock prices. For more information, see an overview of the samples in the Information Center.

How does the server catch clients that disconnect or crash? It creates a subsystem monitor to watch the client and adds an event listener to catch dead events. When a dead event occurs, the server application can clean up the client server state. For example, the server application can remove the client event listener from the dynamic message bean, thereby allowing the server to subscribe to a dynamic topic only when it is needed.

---

## Interoperating with asynchronous beans

The WebSphere Application Server Enterprise Edition asynchronous beans Version 5 service has interoperability problems that affect how the dependent services behave in a clustered environment with mixed WebSphere product versions and mixed platforms. Read the Serialized Asynchronous beans interoperability article and the JMX MBean interoperability issues article to determine whether you need to complete the following general recovery steps.

1. If the Serialized Asynchronous beans interoperability issues apply to you, then apply interim fix PQ72742 or PTF 1 to bring all affected servers up to the appropriate Version `WorkWithExecutionContext` level. If any of the security problems apply to you, then all data must be reserialized. For example, modify the existing Enterprise Java Bean (EJB) or servlet used to serialize the `WorkWithExecutionContext`, so that it rewrites the object using a `java.io.ObjectOutputStream`:

```
ObjectInputStream ois = new ObjectInputStream(...);
ObjectOutputStream oos = new ObjectOutputStream(...);
Object in = ois.readObject();
oos.writeObject(in);
oos.flush();
```

2. If the JMX MBean interoperability issues apply to you:
  - a. Apply interim fix PQ80354 to allow Version 5.0.2 servers to create `ThreadPool` MBeans that are compatible with subsequent versions.
  - b. Change the `mbeanIdentifier` or name fields in all affected scripts and program source to use the new MBean name format rules, if any scripts or programs look up `ThreadPool` MBeans for a `WorkManager`.

## Serialized Asynchronous beans interoperability

Asynchronous beans support taking snapshots of Java 2 Platform Enterprise Edition (J2EE) contextual information from the current servlet or Enterprise Java Beans (EJB) method. This information can include security information, and can be stored using the `WorkManager.create(Work r)` method to get a serializable `WorkWithExecutionContext` object.

If this `WorkWithExecutionContext` object is serialized and stored with customer data, the following information must be reviewed to determine if any steps must be taken to avoid problems.

### Affected products

Both the scheduler and process choreographer services that are included with the WebSphere Application Server Enterprise Edition Version 5 are affected by the problems described in this document. Review the respective sections of each service in the Version 5.0.1 Release Notes, or in one of the following interim fix options for further information:

- PQ72885 - Scheduler Version 5 Interoperability Issues
- PQ72886 - Process Choreographer Version 5 Interoperability Issues



## Asynchronous beans security errors

The asynchronous beans `WorkWithExecutionContext` object created with the 5.0 version of asynchronous beans stores security information incorrectly, which can cause deserialization errors and ASYN9999E errors to appear in the WebSphere Application Server log files. This happens if the interoperating versions of WebSphere Application Server have different implementations of Java Authentication and Authorization Service (JAAS). This error most likely occurs when interoperating between two different platforms, but may occur on the same platform if the JAAS implementation is updated on any of the servers.

Security information is stored with the `WorkWithExecutionContext` object only when the following conditions are true:

- The `WorkManager` resource used to create the `WorkWithExecutionContext` object has the *security* context policy attribute enabled.
- Global Security is enabled on the WebSphere Application Server.

If either of the previous conditions are false, then there should be no interoperability issues.

If the JAAS implementation JAR files have already been updated or changed such that all servers are already experiencing ASYN9999E errors, then apply the original version of the JAAS implementation JAR files that were used to serialize the `WorkWithExecutionContext` data, or follow the Recovery and interoperability steps in the article, *Interoperating with asynchronous beans*, on each platform that is having the problem to ensure all data is reserialized to the appropriate version.

## Asynchronous beans interoperability errors

To avoid further problems interoperating between releases or platforms of asynchronous beans, the serialization mechanisms within asynchronous beans have been updated such so that a 5.0 version of the asynchronous beans object, `WorkWithExecutionContext`, is not able to deserialize a `WorkWithExecutionContext` object that is serialized with a Version 5.0.1 release or later.

## Interoperability and recovery

If any of the conditions described in the *Asynchronous beans security errors* section of this document are true, then all data with the serialized Version 5 `WorkWithExecutionContext` object must be reserialized with the updated format.

Regardless of whether any of the conditions described in the "Asynchronous beans security errors" section of this document are true, apply interim fix, PQ72742, to all WebSphere Application Server Enterprise 5 servers that expect to interoperate with Version 5.0.1 and later servers or expect to interoperate with other Version 5 servers that have PQ72742 applied.

Servers with interim fix, PQ72742, or Version 5.0.1 applied are able to read `WorkWithExecutionContext` data serialized with Version 5, but servers at the Version 5 level are not able to read data serialized in the PQ72742 or Version 5.0.1 format. If this problem occurs, errors such as the following appear: ASYN9999E: Unexpected Exception Occurred: `java.io.OptionalDataException`.

For more information on completing the steps for recovery and interoperability, see the article, *Interoperating with asynchronous beans*.

## Internationalization interoperability issues

There are internationalization interoperability issues with the Asynchronous Beans `WorkWithExecutionContext` object. Read this section and the *Internationalization Interoperability* section of the Version 5.0.1 Release Notes for details on how to determine if this issue affects your work environment, and if so, how to resolve the problems.



Affected customers are those who store Internationalization service context information with the `WorkWithExecutionContext` object. Internationalization information is stored with the `WorkWithExecutionContext` object only when the following conditions are true:

- The `WorkManager` resource used to create the `WorkWithExecutionContext` object has the **com.ibm.ws.i18n** context policy attribute enabled.
- The Internationalization service is enabled on the WebSphere Application Server.

If either of the previous conditions are false on either the application server that *serializes* the `WorkWithExecutionContext` or the application server that *deserializes* the `WorkWithExecutionContext`, no interoperability issue exists.

If both of the previous conditions are true on either of the application servers, then follow the instructions in the Internationalization Interoperability section of the Version 5.0.1 Release Notes or interim fix PQ73371.

## Java Management Extension MBean interoperability

WebSphere Application Server Enterprise Edition Asynchronous Beans service creates several MBeans. Two of these MBeans control the thread pools associated with each configured `WorkManager` object -- the `WorkManager` thread pool and the `AlarmManager` thread pool. Both of these pools can be modified at run time by using a Java Management Extension (JMX) MBean. Each pool has an MBean that can be looked up independently and are named: "`WorkManager:<Workmanager_name>`" and "`<Workmanager_name>:Alarm Pool`" respectively, where, `<WorkManager_name>` is the name field of the configured `WorkManager`.

Because the name for the MBean is not a legal JMX 1.2 `ObjectName`, `AsynchBeans` and the `ThreadPool` manager have been updated to automatically convert all illegal characters to periods ("."). For example, if a `WorkManager` name is **Central WorkManager: Finance,Admin**, the actual MBean name and `mbeanIdentifier` becomes **Central WorkManager. Finance.Admin**.

When looking up a `WorkManager ThreadPool` MBean, all queries for these objects change from:

```
wsadmin>$AdminControl queryNames "WebSphere:*,type=ThreadPool,name=WorkManager:Central WorkManager: Finance,Admin"
```

or

```
wsadmin>$AdminControl queryNames "WebSphere:*,mbeanIdentifier=WorkManager:Central WorkManager: Finance,Admin"
```

to

```
wsadmin>$AdminControl queryNames "WebSphere:*,type=ThreadPool,name=WorkManager.Central WorkManager. Finance.Admin"
```

or

```
wsadmin>$AdminControl queryNames "WebSphere:*,mbeanIdentifier=WorkManager.Central WorkManager. Finance.Admin"
```

---

## Chapter 19. Using object pools

An object pool helps an application avoid creating new Java objects repeatedly. Most objects can be created once, used and then reused. An object pool supports pooling, while waiting to be reused. These object pools are not meant to be used for pooling Java Database Connectivity (JDBC) connections or Java Messaging Service (JMS) connections and sessions. WebSphere Application Server provides specialized mechanisms for dealing with those types of objects. These object pools are intended for pooling application-defined objects or basic Developer Kit types.

To use an object pool, the product administrator must define an *object pool manager* using the administrative console. Multiple object pool managers can be created in an Application Server cell.

**Note:** The Object pool manager service is only supported from within the EJB container or Web container. Looking up and using a configured object pool manager from a Java 2 Platform Enterprise Edition (J2EE) application client container is not supported.

1. Start the administrative console.
2. Click **Resources > Object Pools**.
3. Define the name of the object pool manager. This name can be up to 30 ASCII characters long.
4. Assign the object pool manager a Java Naming and Directory Interface (JNDI) name.
5. Provide a description of this object pool manager.
6. Categorize the object pool manager.

After completing this steps, applications can find the object pool manager by doing a JNDI lookup using the specified JNDI name.

The following code illustrates how an application can find an object pool manager object:

```
InitialContext ic = new InitialContext();
ObjectPoolManager opm = (ObjectPoolManager)ic.lookup("java:comp/env/pool");
```

When the application has an ObjectPoolManager, it can cache an object pool for classes of the types it wants to use. The following is an example:

```
ObjectPool arrayListPool = null;
ObjectPool vectorPool = null;
try
{
 arrayListPool = opm.getPool(ArrayList.class);
 vectorPool = opm.getPool(Vector.class);
}
catch(InstantiationException e)
{
 // problem creating pool
}
catch(IllegalAccessException e)
{
 // problem creating pool
}
```

When the application has the pools, the application can use them as in the following example:

```
ArrayList list = null;
try
{
 list = (ArrayList)arrayListPool.getObject();
 list.clear(); // just in case
 for(int i = 0; i < 10; ++i)
 {
 list.add("" + i);
 }
}
```

```

// do what ever we need with the ArrayList
}
finally
{
 if(list != null) arrayListPool.returnObject(list);
}

```

This example presents the basic pattern for using object pooling. If the application does not return the object, then the only adverse effect is that the object cannot be reused.

---

## Object pool managers

Object pool managers control the reuse of application objects and Developer Kit objects, such as Vectors and HashMaps.

Multiple object pool managers can be created in an Application Server cell. Each object pool manager has a unique cell-wide Java Naming and Directory Interface (JNDI) name. Applications can find a specific object pool manager by doing a JNDI lookup using the specific JNDI name.

The Object pool manager and its associated objects implement the following interfaces:

```

public interface ObjectPoolManager
{
 ObjectPool getPool(Class aClass)
 throws InstantiationException, IllegalAccessException;
 ObjectPool createFastPool(Class aClass)
 throws InstantiationException, IllegalAccessException;
}

public interface ObjectPool
{
 Object getObject();
 void returnObject(Object o);
}

```

Each object pool manager can be used to pool any Java object with the following characteristics:

- The object must be a public class with a public default constructor.
- Each pooled object class must have its own object pool.
- An application gets an object pool for a specific object using either the `ObjectPoolManager.getPool()` method or the `ObjectPoolManager.createFastPool()` method. The difference between these methods is that the `getPool()` method returns a pool that can be shared across multiple threads. The `createFastPool()` method returns a pool that can only be used by a single thread.

If in a Java virtual machine (JVM), the `getPool()` method is called multiple times for a single class, the same pool is returned. A new pool is returned for each call when the `createFastPool()` method is called. Basically, the `getPool()` method returns a pool that is thread-synchronized.

The pool for use by multiple threads is slightly slower than a fast pool because of the need to handle thread synchronization. However, extreme care must be taken when using a fast pool. Consider the following interface:

```

public interface PoolableObject
{
 void init();
 void returned();
}

```

If the objects placed in the pool implement this interface and the `ObjectPool.getObject()` method is called, the object returned has the `init()` method called on it. When the `ObjectPool.returnObject()` method is called, the returned method is called on the object before it is returned to the object pool. Using this method objects can be pre-initialized or cleaned up.

It is not always possible for an object to implement `PoolableObject`. For example, an application might want to pool `ArrayList` objects. The `ArrayList` object needs clearing each time the application reuses it. The application might extend the `ArrayList` object and have the `ArrayList` object implement a poolable object. For example, consider the following:

```
public class PooledArrayList extends ArrayList implements PoolableObject
{
 public PooledArrayList()
 {
 }

 public void init() {
 }

 public void returned()
 {
 clear();
 }
}
```

If the application uses this object, in place of a true `ArrayList` object, the `ArrayList` object is cleared automatically when it is returned to the pool.

Clearing an `ArrayList` object simply marks it as empty and the array backing the `ArrayList` object is not freed. Therefore, as the application reuses the `ArrayList`, the backing array expands until it is big enough for all of the application requirements. When this point is reached, the application stops allocating and copying new backing arrays and achieves the best performance.

It might not be possible or desirable to use the previous procedure. An alternative is to implement a custom object pool and register this pool with the object pool manager as the pool to use for classes of that type. The class is registered by the WebSphere administrator when the object pool manager is defined in the cell. Take care that these classes are packaged in Java Archive (JAR) files available on all of the nodes in the cell where they might be used.

---

## Object pool manager collection

Use this page to manage object pool managers.

To view this administrative console page, click **Resources > Object Pools**.

### Name

The name by which the object pool manager is known for administrative purposes.

<b>Data type</b>	String
<b>Range</b>	1 through 30 ASCII characters

### JNDI Name

The Java Naming and Directory Interface (JNDI) name for the object pool manager.

<b>Data type</b>	String
------------------	--------

## Description

A description of the object pool manager.

**Data type** String

## Category

A category string used to classify or group this object pool manager.

**Data type** String

## Object pool manager settings

Use this page to modify object pool manager settings.

To view this administrative console page, click **Resources > Object Pools > *objectpoolmanager\_name***

### Name

The name by which the object pool manager is known for administrative purposes.

**Data type** String  
**Range** 1 through 30 ASCII characters

### JNDI Name

The Java Naming and Directory Interface (JNDI) name for the object pool manager.

**Data type** String

### Description

A description of the object pool manager.

**Data type** String

### Category

A category string used to classify or to group this object pool manager.

**Data type** String

## Custom object pool collection

Use this page to manage object pools.

To view this administrative console page, click **Resources > Object Pools > *objectpoolmanager\_name* > Object Pools**.

### ***Pool Class Name:***

The fully qualified class name of the objects that are stored in the object pool.

**Data type** String

### ***Pool Impl Class Name:***

The fully qualified class name of the CustomObjectPool implementation class for this object pool.

**Data type** String

### **Custom object pool settings:**

Use this page to modify custom object pool settings.

An object pool manages a pool of arbitrary objects.

To view this administrative console page, click **Resources > Object Pools > objectpoolmanager\_name > Object Pools > objectpool\_name**.

#### *Pool Class Name:*

The fully qualified class name of the objects that are stored in the object pool.

**Data type** String

#### *Pool Impl Class Name:*

The fully qualified class name of the CustomObjectPool implementation class for this object pool.

**Data type** String

---

## **Object pool service settings**

Use this page to enable or disable the object pool service, which manages object pool resources used by the server.

To view this administrative console page, click **Servers > Application Servers > server\_name > Object Pool Service**.

## **Startup**

Specifies whether the server attempts to start the object pool service.

**Default**  
**Range**

Selected  
**Selected**

When the application server starts, it attempts to start the object pool service automatically.

**Cleared**

The server does not try to start the object pool service. If object pool resources are used on this server, then the system administrator must start the object pool service manually or select this property, and then restart the server.

---

## **Object pools: Resources for learning**

Use the following links to find relevant supplemental information about object pools. The information resides on IBM and non-IBM Internet sites, whose sponsors control the technical accuracy of the information.

These links are provided for convenience. Often, the information is not specific to the IBM WebSphere Application Server product, but is useful all or in part for understanding the product. When possible, links

are provided to technical papers and Redbooks that supplement the broad coverage of the release documentation with in-depth examinations of particular product areas.

Furthermore, these links provide guidance on using object pools. Since object pooling is a general topic and the WebSphere Application Server product implementation is only one way to use it, you must understand when object pooling is necessary. These articles help you make that decision.

### **Programming model and decisions**

- Java performance programming, Part 1: Smart object-management saves the day
- Build your own ObjectPool in Java to boost app speed
- Improve the robustness and performance of your ObjectPool
- Java Tip 78: Recycle broken objects in resource pools

---

## **Object pool performance considerations**

Object pools are designed to pool and reuse instances of objects that are both complex and frequently instantiated during run time. To use an object pool, you must define an object pool manager using the administrative console. The object pool manager is available in the Java Naming Directory Interface (JNDI) and can provide object pool instances for any class implementing the PoolableObject interface. The PoolableObject interface defines methods for releasing resources when an object is to be pooled and initializes the internal state when an object is retrieved from the pool. The objects are obtained by methods available on the ObjectPool object.

You should use an object pool when objects are needed repeatedly that are expensive to instantiate. This expense can be gauged through the complexity of the object, as well as the size of the data inside the object. Another consideration is the viability of efficiently cleaning out an object. If the members of an object are large or complex, but cannot be quickly restored to defaults for future use, pooling these objects proves effective. The key to seeing a benefit from using an object pool is to make sure that the total time to retrieve the object from the pool, initialize the internal members, and clear the object for reuse and return it to the pool, is less than the time it takes to instantiate the object, initialize internal members, and subsequently clean up the object when heap space runs out. Retrieving an object from the pool and returning it to the pool have been streamlined as much as possible, but the tasks of member initialization and clearing must fall to the developer to implement and evaluate.

An important issue to consider is when implementing an object pool thread contention. There are two types of object pools to choose from: one is thread-safe and other is not. The latter is not synchronized, allowing for faster access and reduction in contention, but is not recommended in a multithreaded environment. If there is a chance that two threads are operating on a pool at the same time, use the thread-safe object pool. However, be sure to consider how many threads are accessing the pool at the same time and how frequently. If it is more than a few on a regular basis, the single point of contention created in this scenario can sometimes cancel out all benefits with an otherwise well conceived poolable object. A thread-safe pool accessed by multiple threads creates a bigger pool of ready object instances, but lining up to retrieve these instances through synchronization can decrease performance.

When attempting to maximize performance with an object pool, cache the ObjectPoolManager instance wherever it is accessed. By looking up an object once and keeping a reference to it for the future, you can reduce the amount of time spent looking it up in JNDI, which, because of its remote nature, tends to be costly.

Another performance consideration is in the implementation of a poolable returned() method object. Preserve the member to be cleared, if it is needed for next use. This reduces even more clean up and avoids the re-initialization of this member upon reuse. Also, if a member has a more efficient way to clear its state, other than simply destroying it, try to utilize the poolable object returned() method. For example, calling the clear() method on a member Vector is much more efficient than destroying the Vector and instantiating a new one upon reuse.



Using an object pool significantly reduces the amount of memory that applications use and the amount of clean up. Through an intelligent choice of pooled objects and pool types, performance improvements can be achieved by avoiding unnecessary object allocation and clean up.



---

## Chapter 20. Using startup beans

A startup bean is a stateful session bean that is loaded when an application starts. Startup beans enable Java 2 Platform Enterprise Edition (J2EE) applications to execute business logic automatically, whenever an application starts or stops normally.

Startup beans are especially useful when used in with asynchronous bean features. For example, a startup bean might create an alarm object that uses Java Messaging Service (JMS) to periodically publish heartbeat messages on a well-known topic. This enables clients or other server applications to determine whether the application is available.

1. Use the home interface, `com.ibm.websphere.startupservice.AppStartUpHome`, to designate a bean as a startup bean
2. Use the remote interface, `com.ibm.websphere.startupservice.AppStartUp`, to define a `start()` and `stop()` method on the bean.

The startup bean `start()` method is called when the application starts. The startup bean implements any business logic that must run at application start time.

The `start()` method returns a boolean. `True` indicates normal application startup and `false` indicates that the application start process should be aborted. The `start()` and `stop()` methods should not use a `TX_MANDATORY` attribute because there never is a transaction on the thread when the `start()` or `stop()` methods are invoked. Any other `TX_*` attribute can be used. If `TX_MANDATORY` is used, an exception is logged (need a transaction for mandatory) and the application does not start.

The startup bean `stop()` method is called when the application stops and implements any business logic that needs to run. Any exception thrown by a `stop()` method is ignored, but logged to trace.

The `start()` and `stop()` methods on the remote interface use **Run-As** mode. **Run-As** mode specifies the credential information to be used by the security service to determine the permissions that a principal has on various resources. If security is on, the **Run-As** mode needs to be defined on all of the methods called. The identity of the bean without this setting is undefined.

There are no restrictions on what code the `start()` and `stop()` methods can run, since the full Enterprise Application Server programming model is available to these methods.

3. Use an *optional* environment property integer, `wasStartupPriority`, to specify the start order of multiple startup beans in the same Java Archive (JAR) file. If the environment property is found and is the wrong type, application startup is aborted. If no priority value is specified, a default priority of 0 is used. It is recommended that you specify the priority property. Beans that have specified a priority are sorted using this property. Beans with numerically lower priorities are executed first. Beans that have the same priority are executed in an undefined order. All priorities must be positive integers. The priority is used to order beans within an EJB JAR file. The order in which this process happens is applied to different EJB JAR files in a single EAR file is undefined. Beans are stopped in the opposite order to their start priority.



---

## Chapter 21. Using the scheduler service

The scheduler service enables tasks to execute at a requested time. You can schedule the following tasks:

- Invoke a session bean method.
- Send a Java Message Service (JMS) message on a queue or topic.

The scheduler service performs the task, repeating as necessary, according to the metadata for that task.

1. Manage the scheduler service. This article includes instructions for creating and configuring a database for the scheduler, configuring a scheduler instance, and enabling or disabling the scheduler service (the service is enabled by default).
2. Develop and schedule tasks. This article includes instructions for developing various types of tasks, receiving notifications from a scheduler, submitting tasks to a scheduler, and managing tasks.

**Note:** Creating and manipulating scheduled tasks through the Scheduler interface is only supported from within the Enterprise Java Beans (EJB) container or Web container (EJB containers or servlets). Looking up and using a configured scheduler from a Java 2 Platform Enterprise Edition (J2EE) application client container is not supported.

3. Interoperate with the scheduler service. This article explains how to manage the Scheduler service in a clustered environment with mixed WebSphere Application Server product versions and mixed platforms.

---

### Managing the scheduler service

Schedulers are configured using the administrative console and are available to all servers on which the scheduler service is enabled.

1. Create the database for a scheduler.
2. Configure a scheduler.
3. Enable the scheduler service.

The scheduler service is enabled by default.

### Creating the database for a scheduler

Your database system must be installed and available.

It is important to realize that the scheduler uses this database for storing tasks and then executing them. The performance of the scheduler is ultimately limited by the performance of the database. If you need more tasks per second, you can run the scheduler daemons on larger systems or you can use clusters for the session beans used by the tasks. Eventually, however, the task database becomes saturated and you then need a larger or better-tuned database system.

Multiple applications can share a scheduler database. This sharing can lower the cost of administering the scheduler database.

The scheduler requires a database, a Java Database Connectivity (JDBC) provider and a data source.

1. Create the database according to the description for your database system:
  - Creating a Cloudscape database for scheduler.
  - Creating a DB2 database for scheduler.
  - Creating a DB2 database for z/OS for scheduler.
  - Creating an Informix database for scheduler.
  - Creating a Microsoft SQL Server database for scheduler.

- Creating an Oracle database for scheduler.
  - Creating a Sybase 12.0 database for scheduler.
2. If the database is not on the same machine as your IBM WebSphere Application Server, verify that you can access the database from your application server machine.
  3. Configure your JDBC provider and data source. For details, see the topic "Creating and configuring a JDBC provider and data source" in the Information Center.

## Creating a Cloudscape database for scheduler

Cloudscape is a database system implemented in Java code. It is delivered with WebSphere Application Server as three Java Archive (JAR) files. The Cloudscape license that comes with WebSphere Application Server is only for development and testing purposes and not for production purposes.

1. Open a command-line window.
2. Make sure that you have administrator rights for the database system.
3. If you want to use an existing database, skip to step 5. Make sure that the database supports Unicode (UTF-8). Otherwise, it cannot store all characters that can be handled in Java code, and you can run into code page conversion problems when a client uses an incompatible code page.
4. Use the Cloudview utility supplied with the Cloudscape system to create the database, scheddb. Cloudscape supports only one local connection. If WebSphere Application Server is running and accessing a Cloudscape database, then attempts to open a second connection to the database from the command line are rejected.
5. Create the schema.
  - a. Using a text editor, edit the script %WAS\_HOME%\Scheduler\createSchemaCloudscape.ddl according to the instruction at the top of the file.
  - b. Enter one of the following commands (shown on multiple lines for publication): On Windows:

```
%install_root%\java\jre\bin\java -Djava.ext.dirs=%install_root%/lib
-Dij.protocol=jdbc:db2j: -Dij.database=scheddb com.ibm.db2j.tools.ij
%install_root%\Scheduler\createSchemaCloudscape.ddl
```

On UNIX:

```
%install_root%/java/jre/bin/java -Djava.ext.dirs=%WAS_HOME%/lib
-Dij.protocol=jdbc:db2j: -Dij.database=scheddb com.ibm.db2j.tools.ij
%WAS_HOME%/Scheduler/createSchemaCloudscape.ddl
```

The Cloudscape database for the scheduler service exists.

## Creating a DB2 database for scheduler

1. Open a DB2 command-line window.
2. Make sure that you have administrator rights for the database system.
3. If you want to use an existing database, skip to Step 5.

**Note:** Verify that the database supports Unicode (UTF-8). Otherwise, it cannot store all the characters that can be handled in Java code, and you could run into codepage conversion problems when a client uses an incompatible codepage.

To avoid deadlocks, be sure that the DB2 isolation level is set to "read stability". If necessary, enter the command db2set DB2\_RR\_TO\_RS=YES then restart the DB2 instance to activate the change.

4. In the DB2 command line processor, enter this command to create the database:

```
db2 CREATE DATABASE scheddb USING CODESET UTF-8 TERRITORY en-us
```

A DB2 database named scheddb has been created.

5. Create the table space and schema.

- a. Analyze the results of your experiences during development and system testing. The size of your database depends on many factors. If possible, distribute tablespace containers across different logical disks, and implement an appropriate security policy. Consider the performance implications of your choices for buffer pools and log file settings.
- b. Using a text editor, edit the following scripts according to the instruction at the top of each file: %WAS\_HOME%\Scheduler\createTablespaceDB2.dd1, %WAS\_HOME%\Scheduler\createSchemaDB2.dd1, %WAS\_HOME%\Scheduler\dropSchemaDB2.dd1, and %WAS\_HOME%\Scheduler\dropTablespaceDB2.dd1.
- c. Make sure that you are attached to the correct instance. Check the environment variable DB2INSTANCE.
- d. To connect to a database named scheddb, enter the command:

```
db2 connect to scheddb
```

- e. To create the table space, enter the command:

```
db2 -tf createTablespaceDB2.dd1
```

Make sure that the script's output contains no errors. If there were any errors, you can drop the table space using the script dropTablespaceDB2.dd1.

- f. To create the schema (tables and indices), in the DB2 command line processor, enter the command:

```
db2 -tf createSchemaDB2.dd1
```

Make sure that the script's output contains no errors. If there were any errors, you can use the dropSchemaDB2.dd1 file to drop the schema.

The DB2 database for the scheduler exists.

## Creating a DB2 for z/OS database for scheduler

This topic describes how to create a DB2 for z/OS database for the scheduler and how to verify that it is reachable from the application server machine.

1. You must have already installed WebSphere Application Server WebSphere Business Integration Server Foundation on a UNIX or Windows machine.
2. On the z/OS machine that hosts the database:
  - a. Log on to the native z/OS environment.
  - b. If multiple DB2 systems are installed, then decide which subsystem you want to use.
  - c. Make a note of the Internet Protocol (IP) port to which the DB2 subsystem is listening.
  - d. Using the DB2 administration menu, create a new database named SCHEDDB, for example. Note the database name.
  - e. Create a storage group and note the name.
  - f. Decide which user ID is used to connect to the database from the remote machine running the product. Normally, for security reasons, this user ID is not the one you used to create the database.
  - g. Grant the user ID the rights to access the database and storage group. The user ID must also have permission to create new tables for the database.
3. On the server machine:
  - a. Make sure that you have DB2 Connect Gateway (Version 8.1 fix pack 3 or higher) installed. This component is part of the DB2 UDB ESE package, however, you can also install it separately.
  - b. Catalog the remote database using the following commands, either in a script or in a DB2 command line window:



```
catalog tcpip node zosnode remote hostname server IP_port ostype mvs;
catalog database subsystem as subsystem at node zosnode authentication dcs;
catalog dcs database subsystem as subsystem parms ',,INTERRUPT_ENABLED'
```

An important difference exists between DB2 UDB and DB2 for z/OS. DB2 UDB does not have the concept of a subsystem, but DB2 for z/OS does have subsystems. To avoid confusion between Database name and Subsystem name, it is important to understand that because DB2 for z/OS runs in a subsystem, the catalog node and catalog database commands must identify the appropriate subsystem. On DB2 UDB, the subsystem name is not a known concept, and the database name to which it connects is actually the name of the DB2 for z/OS subsystem.

- c. Verify that you can establish a connection to the remote subsystem by entering the following command: `db2 connect to subsystem user userid using password`
- d. Change to the *Scheduler* subdirectory in the application server installation root directory.
- e. Edit the `createTablespaceDB2ZOS.dd1` script. Replace `@STG@` with the storage group name. Replace `@DBNAME@` with the database name (not the subsystem name), and replace `@SCHED_TABLESPACE@` with the name of a valid tablespace.
- f. Run your customized version of `createTablespaceDB2ZOS.dd1`, as described in the header of the script. If this script does not work, or if you want to remove the tablespace, edit and run the `dropTablespaceDB2ZOS.dd1` script.
- g. Edit the `createSchemaDB2ZOS.dd1` script. Replace `@STG@` with the storage group name. Replace `@DBNAME@` with the database name (not the subsystem name). Replace `@TABLE_PREFIX@` with the Table Prefix in the configured scheduler resource, and replace `@SCHED_TABLESPACE@` with a valid tablespace that was created by the `createTablespaceDB2ZOS.dd1` script.
- h. Run your customized version of the `createSchemaDB2ZOS.dd1` script, as described in the header of the script. If this script does not work, or if you want to remove the tables and views, use `dropSchemaDB2ZOS.dd1` to drop the schema.
- i. To avoid deadlocks, verify that the `DB2_RR_TO_RS` DB2 flag is set to **YES**. If necessary, restart the DB2 instance to activate the change.

The DB2 for z/OS database for the scheduler service is created.

## Creating an Informix database for scheduler

1. Open a command-line window.
2. Make sure that you have administrator rights for the database system.
3. If you want to use an existing database, skip to step 5.

**Note:** Make sure that the database supports Unicode (UTF-8) . Otherwise, it cannot store all characters that can be handled in the Java code, and you could run into codepage conversion problems when a client uses an incompatible codepage.

4. If you want to create a new database named `scheddb`, enter the command:

```
dbaccess CREATE DATABASE scheddb with log
```

5. Create the schema.
  - a. Using a text editor, edit the script `%WAS_HOME%\Scheduler\createSchemaInformix.sql` according to the instruction at the top of the file.
  - b. Enter the command:

```
dbaccess scheddb createSchemaInformix.sql
```

The Informix database for scheduler exists.

## Creating a Microsoft SQL Server database for scheduler

1. Open a command-line window.

2. Change to the directory where the configuration scripts for scheduler are located. This is the Scheduler subdirectory of the IBM WebSphere Application Server installation directory. On Windows, enter:

```
cd %WAS_HOME%\Scheduler
```

On UNIX, enter:

```
cd $WAS_HOME/Scheduler
```

3. Using a text editor, edit the schema creation script (createSchemaMSSQL.sql for SQL Server 2000 and createSchemaMSSQL7.sql for SQL Server 7), according to the instructions at the beginning of the file.
4. If you want to use an existing database, skip to step 6.

**Note:** Make sure that the database supports Unicode (UTF-8) . Otherwise, it cannot store all characters that can be handled in the Java code, and you could run into codepage conversion problems when a client uses an incompatible codepage.

5. If you want to create a new database named scheddb:
  - a. Make sure that you are using a user ID that has administrator rights for the database system.
  - b. In the **Enterprise Manager**, expand a server group, then expand a server.
  - c. Right-click **Databases**, then click **New Database**.
  - d. Type the name scheddb.
  - e. Modify any default values, as desired, then save.

A Microsoft SQL Server database named scheddb is created.

6. To create the schema:
  - a. Make sure that you have administrator rights for the database system. The user ID you use to create the schema must be the one that you tell WebSphere Application Server to use when accessing the database.
  - b. Run the script to create the schema (tables and views):

```
isql -S <serverName> -U<userid> -P<password> -D<databaseName> -i <script name>
```

The Microsoft SQL Server database for scheduler exists.

## Creating an Oracle database for scheduler

1. Open a command-line window.
2. Make sure that you have administrator rights for the database system.
3. If you want to use an existing database, skip to step 5.

**Note:** Make sure that the database supports Unicode (UTF-8). Otherwise, it cannot store all characters that can be handled in the Java code, and you could run into codepage conversion problems when a client uses an incompatible codepage.

4. Use the Database Configuration Assistant to create a database named scheddb. Make sure that you select the JServer option for the database. It is recommended to use a Unicode codepage when creating the database. The text data you pass to the APIs must be compatible with the selected codepage.
5. Create the table space and schema.
  - a. Using a text editor, edit the scripts %WAS\_HOME%\Scheduler\createTablespaceOracle.ddl and %WAS\_HOME%\Scheduler\createSchemaOracle.ddl according to the instruction at the top of the files.
  - b. If you do not want the schema to be created in the default instance, set the environment variable ORACLE\_SID
  - c. To create the table space, run the script createTablespaceOracle.ddl. For test purposes you can use the same location for all table spaces and pass the path as a command line argument to the

script, for example, on Windows, user ID scheduser, password schedpwd, database name scheddb, and table space path d:\mydb\ts, enter the command:

```
sqlplus scheduser/schedpwd@scheddb @createTablespaceOracle.ddl d:\mydb\ts
```

If you get any errors creating the table space, you can use dropTablespaceOracle.ddl to drop the table space.

- d. To create the schema, run the script createSchemaOracle.ddl. For example, on Windows systems, enter:

```
sqlplus scheduser/schedpwd@scheddb @createSchemaOracle.ddl
```

If you get any errors creating the schema (tables and views), you can use dropSchemaOracle.ddl to drop the schema.

The Oracle database for scheduler exists.

## Creating a Sybase database for scheduler

1. Open a command-line window.
2. Make sure that you have administrator rights for the database system.
3. Make sure that you have the DTM option for Sybase ASE installed.
4. If you want to use an existing database, skip to step 6.

**Note:** Make sure that the database supports Unicode (UTF-8). Otherwise, it cannot store all characters that can be handled in the Java code, and you could run into codepage conversion problems when a client uses an incompatible codepage.

5. Use the Sybase isql utility to create a database named scheddb. See your Sybase product documentation for details.
6. Create the schema:
  - a. Using a text editor, edit the following script <install\_root>\Scheduler\createSchemaSybase12.ddl according to the instruction at the top of the file.
  - b. Enter the command:

```
isql -S <serverName> -U<userid> -P<password> -D scheddb -i createSchemaSybase12.ddl
```

The Sybase database for the scheduler exists.

## Configuring a scheduler

Before your application can make use of the scheduler service, you need to configure a scheduler instance using the administrative console. Conceptually, a scheduler is similar to a datasource: you specify various configuration attributes, including a JNDI name where the instance will be bound. Once defined, an application using the scheduler API can look up the scheduler object and call various methods to manage tasks.

1. Start the administrative console.
2. Select **Resources > Scheduler Configurations**.
3. Click **New**.
4. Specify configuration settings. Fields marked with an asterisk (\*) are required. The settings are described in detail in the topic Scheduler configuration settings.

## Scheduler configuration collection

Use this page to manage scheduler configurations.

To view this administrative console page, click **Resources > Scheduler Configurations** .

### **Name:**

The name by which this scheduler is known for administrative purposes.

**Data type** String

### **JNDI Name:**

The JNDI name for the scheduler.

The JNDI name specifies where this scheduler instance is bound in the namespace. Clients can look this name up directly, although the use of resource references is recommended.

**Data type** String

### **Description:**

A description of this scheduler for administrative purposes.

**Data type** String

### **Category:**

A string that can be used to classify or group this scheduler.

**Data type** String

### **Datasource JNDI Name:**

Datasource where persistent tasks will be stored.

Any datasource available in the name space can be used with a scheduler. Multiple schedulers can share a single datasource while using different tables by specifying a table prefix.

**Data type** String

### **Datasource Alias:**

Alias to a user name and password used to access the datasource.

**Data type** String

### **Table Prefix:**

String prepended to the table name TASK.

Multiple independent schedulers can share the same database if each instance specifies a different prefix string.

**Data type** String

***Poll Interval:***

The interval at which the scheduler daemon polls the database. Each scheduled, repeating task's minimum repeat interval will be equal to this value regardless of what is specified on the task.

Each poll operation can be expensive. If the interval is extremely small and there are many scheduled tasks, polling can consume a large portion of system resources.

<b>Data type</b>	Integer
<b>Units</b>	Seconds
<b>Default</b>	30
<b>Range</b>	Any positive long integer

***Work Manager:***

Specifies the work manager used by this scheduler.

The Work Manager is a server object that serves as a logical thread pool for the scheduler. Each repeating task that is created using this scheduler will use the "Number Of Alarm Threads" specified in the Work Manager which will affect the number tasks that can run concurrently. Use the Work Manager's "Service Names" property to limit the amount of context information that is propagated to the task when it executes.

When a task fires, the task is run in the Work Manager associated with the scheduler instance. Configuring a scheduler with a specific Work Manager enables you to control how many tasks are actively running at a given time.

***Scheduler configuration settings:***

Use this page to modify scheduler settings.

To view this administrative console page, click **Resources > Scheduler Configurations > scheduler\_name**.

*Name:*

The name by which this scheduler is known for administrative purposes.

**Data type** String

*JNDI Name:*

The JNDI name for the scheduler.

The JNDI name specifies where this scheduler instance is bound in the namespace. Clients can look this name up directly, although the use of resource references is recommended.

**Data type** String

*Description:*

A description of this scheduler for administrative purposes.

**Data type** String

*Category:*

A string that can be used to classify or group this scheduler.

**Data type** String

*Datasource JNDI Name:*

Datasource where persistent tasks will be stored.

Any datasource available in the name space can be used with a scheduler. Multiple schedulers can share a single datasource while using different tables by specifying a table prefix.

**Data type** String

*Datasource Alias:*

Alias to a user name and password used to access the datasource.

**Data type** String

*Table Prefix:*

String prepended to the table name TASK.

Multiple independent schedulers can share the same database if each instance specifies a different prefix string.

**Data type** String

*Poll Interval:*

The interval at which the scheduler daemon polls the database. Each scheduled, repeating task's minimum repeat interval will be equal to this value regardless of what is specified on the task.

Each poll operation can be expensive. If the interval is extremely small and there are many scheduled tasks, polling can consume a large portion of system resources.

<b>Data type</b>	Integer
<b>Units</b>	Seconds
<b>Default</b>	30
<b>Range</b>	Any positive long integer

*Work Manager:*

Specifies the work manager used by this scheduler.

The Work Manager is a server object that serves as a logical thread pool for the scheduler. Each repeating task that is created using this scheduler will use the "Number Of Alarm Threads" specified in the

Work Manager which will affect the number tasks that can run concurrently. Use the Work Manager's "Service Names" property to limit the amount of context information that is propagated to the task when it executes.

When a task fires, the task is run in the Work Manager associated with the scheduler instance. Configuring a scheduler with a specific Work Manager enables you to control how many tasks are actively running at a given time.

## Creating a scheduler resource reference

When a scheduler has been defined in the server configuration, the object instance is bound into the global name space under the configured JNDI name. A resource reference can be used to avoid hardcoding this JNDI name into your application.

You can alternatively create a scheduler resource reference by editing the XML directly. A scheduler resource reference is a (Java 2 Platform Enterprise Edition (J2EE) compliant resource that uses the class `com.ibm.websphere.scheduler.Scheduler` as the object type. For information regarding the XML file format, see the J2EE Specification.

1. Start the Application Assembly Tool Assembly Toolkit.
2. Select your application.
3. In the left-hand panel, right-click on **Resource References** and select **New**.
4. On the **General** tab, complete the following fields:
  - Name** Specify the name suffix. For example, if the scheduler name is *MyScheduler*, the reference JNDI name is **`java:comp/env/MyScheduler`**
  - Type** From the drop-down list select **`com.ibm.websphere.scheduler.Scheduler`**
5. (Optional) You can also specify the global JNDI name to which this resource reference is bound by entering the JNDI name on the **Bindings** tab.

## Scheduler daemon

A scheduler daemon is a background thread that searches for events in the persistent store.

A scheduler daemon is started for each scheduler defined on each server. If "Scheduler 1" is configured on server1, then there will only be one scheduler daemon running on server1 unless it is cloned. If "Scheduler 1" is defined at the node scope level, then the scheduler will run on each server within that node.

The poll interval determines the frequency at which the persistent store is queried. By default, this value is set to 30 seconds. When a task is found that is scheduled to fire within the current poll interval, an alarm is set. The task then runs as close to this time as possible using an alarm thread from the scheduler's associated work manager. Thus, the number of alarm threads configured on the work manager determines how many concurrent tasks are executed. No tasks are lost. If we reach this limit, then new tasks are simply queued to be executed when an alarm thread becomes available. The actual firing time is dictated by server load and availability of free threads in the alarm thread pool of the associated work manager.

## Scheduler daemons in a cluster

When multiple scheduler daemons are configured to the same table (as is the case in a clustered environment), any of the daemons can find a task and set the timer in its Java virtual machine (JVM). The task is executed in the virtual machine where the timer first fires.

## Enabling the scheduler service

Before an application can make use of the scheduler service, you need to configure a scheduler.



The scheduler service manages all schedulers used by a given server. The scheduler service can be enabled and disabled on a server-by-server basis using the administrative console. The service is enabled by default. If you disable the service on a server, all schedulers configured on that server are no longer available. All lookups fail and all scheduler daemons are inactive.

1. Start the administrative console.
2. Select **Servers** > *server\_name* > **Scheduler Service**.
3. Select or clear the **Startup** check box to enable or disable the service.
4. Click **Save** on the menu bar to save your configuration.

The change takes effect on the next server restart.

## Scheduler service settings

Use this page to enable or disable the scheduler service, which manages scheduler resources used by the server.

To view this administrative console page, click **Servers** > **Application Servers** > *server\_name* > **Scheduler Service**.

### Startup:

Specifies whether the server will attempt to start the scheduler service.

**Default**  
**Range**

Selected

**Selected**

When the application server starts, it attempts to start the scheduler service automatically.

**Cleared**

The server does not try to start the scheduler service. If scheduler resources are to be used on this server, the system administrator must start the scheduler service manually or select this property, then restart the server.

---

## Developing and scheduling tasks

1. Developing a task.

The scheduler API supports different implementations of the TaskInfo interface, each of which can be used to schedule a particular type of work. Refer to one of the following topics for details:

- Developing a task that calls a session bean.
- Develop a task that sends a Java Messaging Service (JMS) message. This task object can send a JMS message to either a queue or a topic.

**Note:** Create and manipulate scheduled tasks through the Scheduler interface is only supported from within the EJB container or Web container (enterprise beans or servlets). Looking up and using a configured scheduler from a J2EE application client container is not supported.

2. Receiving scheduler notifications. A notification sink is set on a task in order to receive the notification events that are generated by a scheduler when it performs an operation on the task.
3. Submitting a task to a scheduler. After a TaskInfo object has been created, it can be submitted to the scheduler for task creation by calling the Scheduler.create() method.
4. Managing tasks with a scheduler.

## Developing a task that calls a session bean

The scheduler API supports different implementations of the `TaskInfo` interface, each of which can be used to schedule a particular type of work. This topic describes how to call a method on a task handler session bean using the `BeanTaskInfo` implementation.

1. Create a stateless session bean that implements the `process()` method in the `com.ibm.websphere.scheduler.TaskHandler` remote interface. Place the business logic you want created in the `process()` method. The `process()` method is called when the task fires. The `Home` and `Remote` interfaces must be set as follows in the bean's deployment descriptor:
  - `com.ibm.websphere.scheduler.TaskHandlerHome`
  - `com.ibm.websphere.scheduler.TaskHandler`
2. In a java server page (JSP), servlet or session bean, create an instance of the `BeanTaskInfo` interface by using the following `Scheduler.createBeanTaskInfo()` method:

```
//lookup the scheduler instance to be used
Scheduler scheduler = (Scheduler)new InitialContext.lookup("java:comp/env/Scheduler");

BeanTaskInfo taskInfo = (BeanTaskInfo) scheduler.createTaskInfo(BeanTaskInfo.class)
```

**Note:** Creating a `BeanTaskInfo` object does not add the task to the persistent store. Rather, it creates a placeholder for the necessary data. The task is not added to the persistent store until the `create()` method is called on a `Scheduler` instance, as described in the topic [Submitting a task to a scheduler](#).

3. Now, set parameters on the `BeanTaskInfo` object. These parameters define which task is to run, which session bean is called, and so on.

The `TaskInfo` interface contains various `set()` methods that you can use to control execution of the task, including when the task will fire and what work the task will do when it fires. For example:

```
//create a date object which represents 30 seconds from now
java.util.Date startDate = new java.util.Date(System.currentTimeMillis()+30000);

//find the session bean to be called when the task executes
Object o = new InitialContext().lookup("java:comp/env/ejb/MyTaskHandlerHome");
TaskHandlerHome home = (TaskHandlerHome)javax.rmi.PortableRemoteObject.narrow(o,TaskHandlerHome.class);

//now set the start time and task handler to be called in the task info
taskInfo.setTaskHandler(home);
taskInfo.setStartTime(startDate);
```

The `TaskInfo` interface specifies additional control points, as documented in Javadoc.

A `TaskInfo` object has been created that contains all of the relevant data for a task.

Submit the task to a scheduler instance for creation, as described in the topic [Submitting a task to a scheduler](#).

## Developing a task that sends a JMS message

The scheduler API supports different implementations of the `TaskInfo` interface, each of which can be used to schedule a particular type of work. This topic describes how to use the `MessageTaskInfo` implementation, which sends a JMS message to either a queue or a topic.

1. Create an instance of the `MessageTaskInfo` interface by using the following `Scheduler.createMessageTaskInfo()` method:

```
//lookup the scheduler instance to be used
Scheduler scheduler = (Scheduler)new InitialContext.lookup("java:comp/env/Scheduler");

MessageTaskInfo taskInfo = (MessageTaskInfo) scheduler.createTaskInfo(MessageTaskInfo.class);
```

**Note:** Creating a `MessageTaskInfo` object does not add the task to the persistent store. Rather, it creates a placeholder for the necessary data. The task is not added to the persistent store until the `create()` method is called on a `Scheduler` instance, as described in the topic `Submitting a task to a scheduler`.

2. Set parameters on the `MessageTaskInfo` object.

The `TaskInfo` interface contains various `set()` methods that can be used to control execution of the task, including when the task will fire and what work the task will do when it fires. For example:

```
//create a date object which represents 30 seconds from now
java.util.Date startDate = new java.util.Date(System.currentTimeMillis()+30000);
```

```
//now set the start time and the JNDI names for the queue connection factory and the queue
taskInfo.setConnectionFactoryJndiName("jms/MyQueueConnectionFactory");
taskInfo.setDestination("jms/MyQueue");
taskInfo.setStartTime(startDate);
```

The `TaskInfo` interface specifies additional control points, as documented in Javadoc.

A `TaskInfo` object has been created that contains all of the relevant data for a task.

Submit the task to a scheduler instance for creation, as described in the topic `Submitting a task to a scheduler`.

## Receiving scheduler notifications

Various notification events are generated by a scheduler when it performs an operation on a task. These events include:

### **Scheduled**

A task has been scheduled.

### **Purged**

A task has been permanently deleted from the persistent store.

### **Suspended**

A task was suspended.

### **Resumed**

A task was resumed.

### **Complete**

A task has run completely. If it was a repeating task, all repeats have been performed.

### **Cancelled**

A task has been cancelled. It will not run again.

**Fired** A task fired successfully.

### **Fire Failed**

A task could not fire successfully.

To receive notification events, call the `setNotificationSink()` method on the `TaskInfo` interface before creating the event. The `setNotificationSink()` method enables you to specify the session bean that is to act as the callback, and a mask that restricts which events are generated.

1. Create a notification sink session bean. Create a stateless session bean that implements the `handleEvent()` method in the `com.ibm.websphere.scheduler.NotificationSink` remote interface. The `handleEvent()` method is called when the notification is fired. The `Home` and `Remote` interfaces can be set as follows in the bean's deployment descriptor:

```
com.ibm.websphere.scheduler.NotificationSinkHome
com.ibm.websphere.scheduler.NotificationSink
```

The notification sink bean must exist in the same application (EAR file) that is used to create the task. The `NotificationSink` interface defines the following method:

```
public void handleEvent(TaskNotificationInfo task) throws java.rmi.RemoteException;
```

The transactional context used by the session bean is defined by the assembler.

2. Specify the notification sink session bean to be used as the callback. The following code illustrates how to set this option:

```
TaskInfo taskInfo = ...
Object o = new InitialContext().lookup("java:comp/env/ejb/NotificationSink");
NotificationSinkHome home =
 (NotificationSinkHome)javax.rmi.PortableRemoteObject.narrow(o,NotificationSinkHome.class);
taskInfo.setNotificationSink(home,TaskNotificationInfo.ALL_EVENTS);
```

3. Specify the event mask. The event mask is specified as an integer mask. You can either use an individual mask such as `TaskNotificationInfo.CREATED` to receive specific events, `TaskNotificationInfo.ALL_EVENTS` to receive all events or a combination of specific events.. For example:

```
int eventMask = TaskNotificationInfo.CREATED+TaskNotificationInfo.PURGED;
```

## Submitting a task to a scheduler

This task assumes that you have already configured a scheduler instance using the administrative console.

Once you have developed a `TaskInfo` object that contains all relevant data for a task, submit the task to a scheduler instance for creation. For example:

```
//lookup the scheduler instance to be used
Scheduler scheduler = (Scheduler)new InitialContext.lookup("java:comp/env/Scheduler");
```

```
TaskStatus status = scheduler.create(taskInfo);
```

When you configure a scheduler, it is bound to a global JNDI name. Although the desired scheduler instance can be found by performing a lookup on that JNDI name, it is better to create a resource reference, which allows for more flexibility in configuring the scheduler.

**5.1 +** The scheduler service utilizes the asynchronous beans deferred start mechanism to propagate Java 2 Enterprise Edition (J2EE) service context information to a task when it runs. The service context information captured when the task is created and later applied to the task is controlled by the Service Context settings on the `WorkManager` configuration object, which the scheduler configuration references. For example, security and internationalization can be enabled. See asynchronous beans for details on how to configure the Application Server to propagate these service contexts.

**5.1 +** To secure your application using the global security mechanism built into the product, create each task with the correct credentials on the thread. When each task has the correct credentials, you can disable and re-enable security without causing any security problems. If you do not set the security context when the scheduler task is created and you later enable security in the target application, then a security exception or error message might display.

**5.1 +** The security context is not set if any of the follow conditions are true:

- Global security is disabled.
- The security context policy is disabled on the configured `WorkManager` for the associated scheduler instance configuration.
- A credential is not set on the thread. For example, enterprise beans or the servlet that is used to create the scheduled task is not secured.

**5.1 +** If any of the previously mentioned conditions are true when you create your task, and you need to enable security on your application server or application, then you must complete the following steps:

1. Cancel the task with security disabled.

2. Recreate each scheduled task with security enabled. Recreating each task enables the scheduler to store the desired security context with the task.

**Note:** The scheduler interface is a local interface. It can only be used by server-side code; that is, J2EE applications.

Once the call to the `create()` method is complete, the task exists in the persistent store and is run at the time specified in the `TaskInfo` object. This call is transactional. If a transactional context is present on the thread when the `create()` method rolls back or is aborted, the task does not run.

The status object, which has been returned by the call to the `create()` method, contains information about the state of the task, as well as the task ID. The task ID is the unique identifier for this task, and is required if the task is to be suspended, resumed, cancelled, and so on, at a later time.

**Note:** The status object is only a snapshot of the current state of the task. Use the `Scheduler.getStatus()` method to receive the current state when needed.

## Managing tasks with a scheduler

When a task is created by calling the `create()` method on a scheduler instance, a `TaskStatus` object is returned to the caller. The status object contains the task ID, which is a unique identifier. The scheduler API defines several additional methods that pertain to the management of tasks, each of which accepts the task ID as a parameter. The following task management methods are defined:

### **suspend()**

Suspends a task. The task does not run until it has been resumed.

### **resume()**

Resumes a previously suspended task.

### **cancel()**

Cancels a task. The task is not run.

### **purge()**

Permanently deletes the task from the persistent store.

### **getStatus()**

Returns the current status of the task.

For example, the following code creates and cancels a task:

```
//Create the task.
TaskInfo taskInfo = ...
TaskStatus status = scheduler.create(taskInfo);

//Get the task ID
String taskId = status.getTaskId();

//Cancel the task. Specify the purgeAlso flag so that the task does not remain in the persistent store
scheduler.cancel(taskId,true);
```

**Transactionality.** All methods of the scheduler API are transactional. If a global transactional context is present, it is used to perform the operation. If an unexpected exception is thrown, the transaction is marked to roll back. If an expected or declared exception is thrown, the transaction remains intact and the caller must choose to roll back or to commit the transaction. If the transaction is rolled back at some point, all scheduler operations performed within the transaction were also rolled back.

If a local transactional context is present, it is suspended and a new global transactional context begins. Likewise, if no transactional context is active, a global transactional context begins. In both cases, if an unexpected exception is thrown, the transaction rolls back. If a declared exception is thrown, the transaction is committed.

If another thread is concurrently modifying the task in question, a `TaskPending` exception is thrown. This is because schedulers lock the database optimistically. The calling application can then retry the operation.

All methods defined by the scheduler API are described in Javadoc.

## Transactions and the scheduler service

### Transactions and the scheduler daemon

Scheduled `BeanTaskInfo` and `MessageTaskInfo` objects are guaranteed to execute only once. This is accomplished by grouping all of the work done in the task as a single unit of work. When each task fires, the following events occur in a single global transactional context:

1. The context of the application that created the task is applied to the thread.
2. A global transactional context is started.
3. The next fire time and start-by time are calculated using the `UserCalendar` bean or the `DefaultUserCalendar`.
4. The task database task record is updated in the database with the state of the next task or deleted if the task is complete and the task's auto-purge setting is true.
5. The `BeanTaskInfo` or `MessageTaskInfo` object is executed.
6. If the task fails and the `NotificationSink` bean is set, a `FIRE_FAILED` notification is fired on a separate transaction.
7. If the task's `NotificationSink` bean is set, then the various notifications are fired as required.
8. The global transaction is committed.

Because all of a task's events are executed in a single global transactional context, you need to consider the following points in order to avoid transaction-related errors:

- Each resource participating in the task's transaction must be 2-phase XA capable.  
This includes the JDBC datasource configured for the scheduler, any JMS services used by the `MessageTaskInfo` objects, and any resources used within any of the `UserCalendar`, `TaskHandler`, or `NotificationSink` beans that have a transaction setting of "Requires".
- One resource can be single-phase, if last participant support is enabled for the application that created the transaction.

All unexpected exceptions are logged to the activity log and all events participating in the task's global transaction are rolled back. This includes changes to the task's database record, which force the task to be executed again when the scheduler daemon polls the database during the next poll cycle. The `UserCalendar`, `TaskHandler`, and `NotificationSink` beans can choose not to participate in the global transaction by setting the bean's transaction setting to "Requires new".

### Transactions and the scheduler interface

All Scheduler interface methods participate in a single global transactional context. If a global transactional context is already present on the thread when the `create()`, `suspend()`, `resume()`, `cancel()`, and `purge()` methods are executed, the existing global transaction is used. Otherwise, a new global transaction begins.

If the method participates in the caller's global transaction and an unexpected error occurs, the transaction is marked to roll back. If the exception is a declared exception, then the exception is rethrown to the caller, and the transaction is left alone for the caller to commit or roll back.

If the method starts its own global transaction and any exception occurs, the transaction is rolled back, and the exception is rethrown to the caller.

## Scheduler interface

A scheduler object exists in the JNDI namespace for each scheduler configuration. A reference to a scheduler can be obtained by performing a lookup on the JNDI name; however, the lookup is valid only

from the server process where the scheduler instance exists. Once a reference has been obtained, tasks can be created, suspended, cancelled, and so on, if the caller has access to the scheduler instance.

For details, see the Interface Scheduler in the Javadoc.

#### **Task creation**

The task is created in the persistent store using the caller's global transactional context if present. See the topic "Transactions and the scheduler service" on page 1110 for more details. Since this is a transactional operation, the task cannot be run or modified from another thread until the current transaction commits.

#### **Task modification**

Tasks that have been created can be modified with the `suspend()`, `resume()`, `cancel()`, and `purge()` methods. These methods take a Task Identifier string as a parameter, which is generated by the `create()` method and can be found in the `TaskStatus` object. If a task is currently running or being modified by another thread, an operation that attempts to modify the state of the task does not block on the attempt, but a `TaskPending` exception is thrown. The operation can be reattempted at another time. Tasks can only be modified by the same application (EAR file) that was used to create the task.

#### **Task execution**

Tasks are executed in the thread pool specified by the configuration's work manager, under the security ID of the task creator. If multiple schedulers are configured to share the same database table, the tasks found in the table can be executed on any of the schedulers, whether or not they are in the same server, node, or cell.

#### **Task lookup**

Tasks can be located using the Name property that was assigned at creation time. This is useful when you need to modify a group of tasks and tracking individual task ID's is not convenient.

### **TaskInfo interface**

`TaskInfo` objects contain the information that can be used to create a task. Several implementations of this class exist, one for each type of task that can be run. Available `TaskInfo` implementations include:

#### **BeanTaskInfo**

Calls a stateless session bean.

#### **MessageTaskInfo**

Sends a JMS message to a queue or publishes a message to a topic.

For details, see the Interface `TaskInfo` in the Javadoc.

After a `TaskInfo` object is created, it can be submitted to the scheduler for task creation by calling the `Scheduler.create()` method.

**Specifying time intervals.** `setStartTimeInterval()`, `setStartByInterval()`, and `setRepeatInterval()` methods all take a `String` parameter that represents time interval. Time intervals are calculated using *user calendars*.

### **TaskHandler interface**

A task handler is a user-defined stateless session bean that is called by tasks created using a `BeanTaskInfo` object. A task handler bean uses the following home and remote interfaces, which are defined in the deployment descriptor using the Application Assembly Tool Assembly Toolkit or WebSphere Studio Application Developer:

```
com.ibm.websphere.scheduler.TaskHandlerHome
com.ibm.websphere.scheduler.TaskHandler
```

The bean itself needs to implement the `process()` method defined in the remote interface. For details, see the Interface `TaskHandler` in the Javadoc.



If a task is created using a `BeanTaskInfo` object, the `process()` method on the `TaskHandler` session bean is called whenever the task runs. Because the `TaskStatus` object for the task is passed as a parameter, the handler can make use of the saved `UserContext` field, as well as determine information about the task, such as when it will fire next, the number of repeats remaining, and so on.

## NotificationSink interface

A notification sink is a user-defined stateless session bean that is called by tasks when their state changes throughout the bean's lifecycle. A notification sink bean uses the following home and remote interfaces, which are defined in the deployment descriptor using the Application Assembly Tool Assembly Toolkit or WebSphere Studio Application Developer::

```
com.ibm.websphere.scheduler.NotificationSinkHome
com.ibm.websphere.scheduler.NotificationSink
```

The bean itself needs to implement the `handleEvent()` method defined in the remote interface. For details, see the Interface `NotificationSink` in the Javadoc.

A `NotificationSink` provides an event notification callback on a task-by-task basis. A notification sink is set on the `TaskInfo` interface, using the `setNotificationSink()` method. If a notification sink is not specified on a task, all notifications are lost; however, the status of a task can be determined by calling the `getStatus()` method from the `Scheduler` interface. A notification callback is made for each of the following events:

- Scheduled
- Suspended
- Resumed
- Fired
- Fire Failed
- Complete
- Purged

## UserCalendar interface

A user calendar is a user-defined stateless session bean that is called by tasks when they need to calculate date-related values. A user calendar bean uses the following home and remote interfaces, which are defined in the deployment descriptor using the Application Assembly Tool Assembly Toolkit or WebSphere Studio Application Developer:

```
com.ibm.websphere.scheduler.UserCalendarHome
com.ibm.websphere.scheduler.UserCalendar
```

The bean itself needs to implement the `applyDelta()` and `validate()` methods defined in the remote interface. For details, see the Interface `UserCalendar` in the Javadoc.

User calendars can be used to calculate time intervals, such as the time between when a repeating task fires and the next time it fires. A user calendar takes a `java.util.Date` object and applies the interval string. The resulting object is a `java.util.Date` object that is an incremented date.

User calendars are set by the `setUserCalendar()` method on the `TaskInfo` interface, and called by the scheduler run-time code when a delta calculation is necessary.

The following methods on the `TaskInfo` interface specify delta strings that use the user calendar for calculation:

- `setStartTimeInterval`
- `setStartByInterval`
- `setRepeatInterval`

### Default user calendar

If a user calendar has not been specified using the `TaskInfo.setUserCalendar()` method, a default

user calendar is used. The default calendar allows for simple delta specifications, such as seconds, minutes, hours, days, and months. See the Javadoc for details on the default calendar.

### Calendar specifiers

A single user calendar can contain logic for multiple calendars. Which calendar is used is determined by a string that acts as the specifier. For example, a bean might be implemented to recognize the interval "day", with a specifier that determines whether to calculate "day" as a standard calendar day, or as a business day.

### Internationalization and time zones

Scheduler makes use of the `java.util.Date` class when storing and processing dates. Internally, this class saves the time as milliseconds since the Epoch, Greenwich Mean Time. Since the `Date` is not converted to local time until converted to a string, the scheduler respects the time zone where the date was created.

### Writing user calendars

Because the user calendar is a stateless session bean, the same Java 2 Platform Enterprise Edition (J2EE) programming model available to other session beans is available to the user calendar as well.

---

## Interoperating with the Scheduler service

The WebSphere Application Server Enterprise Edition Scheduler Version 5.0 service has several interoperability problems that affect how the Scheduler behaves in a clustered environment with mixed WebSphere versions and/or mixed platforms. Specifically, there are two issues that cause interoperability problems:

1. The EJB HomeHandle format has changed between the Versions 5.0 and 5.0.1 of the product. See interim fix PQ72184 or the Interoperability section of the Version 5.0.1 Release Notes.
2. The asynchronous beans serialization mechanism has changed between Versions 5.0 and 5.0.1 of the product. See interim fix PQ72742 or the asynchronous beans section of the Version 5.0.1 Release Notes.

To avoid problems interoperating between Version 5.0 and 5.0.1 and or higher, all Version 5.0 scheduler service customers must first apply the Version 5.0 fix packs associated with interim fixes PQ72184 and PQ72742 and remove all scheduled tasks that were created prior to applying these fix packs.

Be aware that removing all scheduled tasks involves removing data from your scheduler service databases. Follow these procedures for each scheduler instance defined on the affected server.

1. Recreate all scheduler tasks. Use this method if you wish to programmatically recreate your scheduled tasks.
2. Delete all scheduler tasks. Use this method to remove all of your Scheduler tasks and recreate them manually.
3. Recreate all scheduler tables. Use this method to drop and recreate the scheduler tables.

## Recreating Scheduler tasks

This step requires advanced knowledge of developing J2EE Applications and the Scheduler programming interfaces.

1. Use the administrative console to locate each Scheduler resource that was used in WebSphere Application Server Version 5.0. For each Scheduler resource, note the JNDI name.
2. Back up the Scheduler database.
3. Create a new or modify an existing EJB or servlet J2EE application to implement as a method that you intend to use as the Scheduler update program.
4. Create a method similar to the included example. This action finds all existing Scheduler tasks, deletes them and creates new ones with the same parameters. ***If Global Security is enabled in the WebSphere Application Server and the "security" context is enabled on the WorkManager***

*referenced by the Scheduler resource, the current security context is used (as well all other J2EE contexts that are enabled on the creating EJB or servlet thread).* See the following as an example:

```
public void recreateTasks(String schedulerJNDIName)
 throws Exception
{
 InitialContext ctx = new InitialContext();
 Scheduler s = (Scheduler)ctx.lookup(schedulerJNDIName);

 Iterator tasks = null;;
 try
 {
 tasks = s.findTasksByName("%");
 }
 catch (SchedulerNotAvailableException e)
 {
 e.printStackTrace();
 throw e;
 }

 // Iterate through each task and recreate it.
 while(tasks.hasNext())
 {
 TaskInfo curTask = (TaskInfo) tasks.next();

 int retries=0;
 boolean deleted=false;
 TaskStatus status=null;

 // It's best to include each cancel/create
 // in it's own transaction (not shown here).
 while(!deleted && retries < 5)
 {
 try
 {
 // Delete the task.
 s.cancel(curTask.getTaskId(), true);
 deleted = true;

 // Create a new one.
 int createRetries = 0;
 boolean created = false;
 while(!created && createRetries<5)
 {
 try
 {
 s.create(curTask);
 created = true;
 }
 catch (Exception e)
 {
 ++createRetries;
 Thread.sleep(5000);
 }
 }
 }
 }
 }
}
```

```

 if (!created)
 {
 System.out.println("Task Not Created: " + curTask.getTaskId());
 }
 }
 catch (Exception se)
 {
 ++retries;
 Thread.sleep(5000);
 }
}
if (!deleted)
{
 System.out.println("Task Not Deleted: " + curTask.getTaskId());
}
}
}
}

```

## Deleting Scheduler tasks

Use this method to remove *all* of your Scheduler tasks and recreate them manually.

1. Use the administrative console to locate each Scheduler resource that was used in WebSphere Application Server Version 5.0. For each Scheduler resource, note the following:
  - JDBC DataSource JNDI Name
  - Table Prefix
2. Use the administrative console to locate each JDBC DataSource, and note the connection and location parameters.
3. Back up the Scheduler database.
4. Using the native database utilities, remove all of the records for each Scheduler resource for each database noted in steps 1 and 2. The following SQL query template can be used to remove all records. Substitute TBLPFX with the actual Table Prefix.

```
DELETE FROM<TBLPFX>TASK
```

Once the tasks are deleted, new tasks can be created immediately. In addition, the Scheduler service can be active while executing the DELETE SQL command in step 4. However, the operation might take longer to execute or even time out if the Scheduler has database records locked.

## Recreating Scheduler tables

Use this method to drop and recreate the Scheduler tables.

1. Use the administrative console to locate each Scheduler resource that was used in WebSphere Application Server Version 5.0. For each Scheduler resource, note the following:
  - JDBC DataSource JNDI Name
  - Table Prefix
2. Use the administrative console to locate each JDBC DataSource, and note the connection and location parameters.
3. Back up the Scheduler database.
4. For each database and Table Prefix identified in the previous steps, execute the dropSchemaXXX.ddl and createSchemaXXX.ddl scripts located in the <install\_root>/Scheduler directory of your WebSphere Application Server, where <install\_root>, is the directory where the WebSphere Application Server is installed. Each DDL script has instructions (located at the beginning of the

respective script and in the form of a comment) on how to edit and execute the script. Refer to the Creating the database for scheduler article for more information.

---

## Chapter 22. Using shared work areas

The work area service enables application developers to implicitly propagate information beyond the information passed in remote calls. Applications can create a work area, insert information into it, and make remote invocations. The work area is propagated with each remote method invocation, eliminating the need to explicitly include an appropriate argument in the definition of each method. The methods on the server side can use or ignore the information in the work area as appropriate.

Before proceeding with the steps to implement work areas, as described below, review the topic `WorkArea service: Overview`.

1. Developing applications that use work areas. Applications interact with the work area service by implementing the `UserWorkArea` interface.
2. Managing work areas. The work area service is managed using the administrative console.

---

### WorkArea service - Overview

One of the foundations of distributed computing is the ability to pass information, typically in the form of arguments to remote methods, from one process to another. When application-level software is written over middleware services, many of the services rely on information beyond that passed in the application's remote calls. Such services often make use of the implicit propagation of private information in addition to the arguments passed in remote requests; two typical users of such a feature are security and transaction services. Security certificates or transaction contexts are passed without the knowledge or intervention of the user or application developer. The implicit propagation of such information means that application developers do not have to manually pass the information in method invocations, which makes development less error-prone, and the services requiring the information do not have to expose it to application developers. Information such as security credentials can remain secret.

The `WorkArea` service gives application developers a similar facility. Applications can create a work area, insert information into it, and make remote invocations. The work area is propagated with each remote method invocation, eliminating the need to explicitly include an appropriate argument in the definition of every method. The methods on the server side can use or ignore the information in the work area as appropriate. If methods in a server receive a work area from a client and subsequently invoke other remote methods, the work area is transparently propagated with the remote requests. When the creating application is done with the work area, it terminates it.

There are two prime considerations in deciding whether to pass information explicitly as an argument or implicitly by using a work area. These considerations are:

- **Pervasiveness:** Is the information used in a majority of the methods in an application?
- **Size:** Is it reasonable to send the information even when it will not be used?

When information is sufficiently pervasive that it is easiest and most efficient to make it available everywhere, application programmers can use the `WorkArea` service to simplify programming and maintenance of code. The argument does not need to go onto every argument list. It is much easier to put the value into a work area and propagate it automatically. This is especially true for methods that simply pass the value on but do nothing with it. Methods that make no use of the propagated information simply ignore it.

Work areas can hold any kind of information, and they can hold an arbitrary number of individual pieces of data, each stored as a property.

### Work area property modes

The information in a work area consists of a set of properties; a property consists of a key-value-mode triple. The key-value pair represents the information contained in the property; the key is a name by which the associated value is retrieved. The mode determines whether the property can be removed or modified.

## Property modes

There are four possible mode values for properties, as shown in the following code example:

### Code example: The PropertyModeType definition

```
public final class PropertyModeType {
 public static final PropertyModeType normal;
 public static final PropertyModeType read_only;
 public static final PropertyModeType fixed_normal;
 public static final PropertyModeType fixed_readonly;
};
```

A property's mode determines three things:

- Whether the value associated with the key can be modified
- Whether the property can be deleted
- Whether the mode associated with the key-value pair can be modified

The two read-only modes forbid changes to the information in the property; the two fixed modes forbid deletion of the property.

The WorkArea service does not provide methods specifically for the purpose of modifying the value of a key or the mode associated with a property. To change information in a property, applications simply rewrite the information in the property; this has the same effect as updating the information in the property. The mode of a property governs the changes that can be made. Modifying key-value pairs describes the restrictions each mode places on modifying the value and deleting the property. Changing modes describes the restrictions on changing the mode.

## Changing modes

The mode associated with a property can be changed only according to the restrictions of the original mode. The read-only and fixed read-only properties do not permit modification of the value or the mode. The fixed normal and fixed read-only modes do not allow the property to be deleted. This set of restrictions leads to the following permissible ways to change the mode of a property within the lifetime of a work area:

- If the current mode is normal, it can be changed to any of the other three modes: fixed normal, read-only, fixed read-only.
- If the current mode is fixed normal, it can be changed only to fixed read-only.
- If the current mode is read-only, it can be changed only by deleting the property and re-creating it with the desired mode.
- If the current mode is fixed read-only, it cannot be changed.
- If the current mode is not normal, it cannot be changed to normal. If a property is set as fixed normal and then reset as normal, the value is updated but the mode remains fixed normal. If a property is set as fixed normal and then reset as either read-only or fixed read-only, the value is updated and the mode is changed to fixed read-only.

**Note:** The key, value, and mode of any property can be effectively changed by terminating (completing) the work area in which the property was created and creating a new work area. Applications can then insert new properties into the work area. This is not precisely the same as changing the value in the original work area, but some applications can use it as an equivalent mechanism.

## Nested work areas

Applications can nest work areas. When an application creates a work area, a work area context is associated with the creating thread. If the application thread creates another work area, the new work area is nested within the existing work area and becomes the current work area. Nested work areas allow applications to define and scope properties for specific tasks without having to make them available to all



parts of the application. All properties defined in the original, enclosing work area are visible to the nested work area. The application can set additional properties within the nested work area that are not part of the enclosing work area.

An application working with a nested work area does not actually see the nesting of enclosing work areas. The current work area appears as a flat set of properties that includes those from enclosing work areas. In the figure below, the enclosing work area holds several properties and the nested work area holds additional properties. From the outermost work area, the properties set in the nested work area are not visible. From the nested work area, the properties in both work areas are visible.




Figure 25. Defining new properties in nested work areas

Nesting can also affect the apparent settings of the properties. Properties can be deleted from or directly modified only within the work areas in which they were set, but nested work areas can also be used to temporarily override information in the property without having to modify the property. Depending on the modes associated with the properties in the enclosing work area, the modes and the values of keys in the enclosing work area can be overridden within the nested work area.

The mode associated with a property when it is created determines whether nested work areas can override the property. From the perspective of a nested work area, the property modes used in enclosing work areas can be grouped as follows:

- Modes that permit a nested work area to override the mode or the value of a key locally. The modes that permit overriding are:
  - Normal
  - Fixed normal
- Modes that do not permit a nested work area to override the mode or the value of a key locally. The modes that do not permit overriding are:
  - Read-only
  - Fixed read-only

If an enclosing work area defines a property with one of the overridable modes, a nested work area can specify a new value for the key or a new mode for the property. The new value or mode becomes the value or mode seen by subsequently nested work areas. Changes to the mode are governed by the

restrictions described in Changing modes. If an enclosing work area defines a property with one of the modes that cannot be overridden, no nested work area can specify a new value for the key.

A nested work area can delete properties from enclosing work areas, but the changes persist only for the duration of the nested work area. When the nested work area is completed, any properties that were added in the nested area vanish and any properties that were deleted from the nested area are restored.

The following figure illustrates the overriding of properties from an enclosing work area. The nested work area redefines two of the properties set in the enclosing work area. The other two cannot be overridden. The nested work area also defines two new properties. From the outermost work area, the properties set or redefined in the nested work are not visible. From the nested work area, the properties in both work areas are visible, but the values seen for the redefined properties are those set in the nested work area.

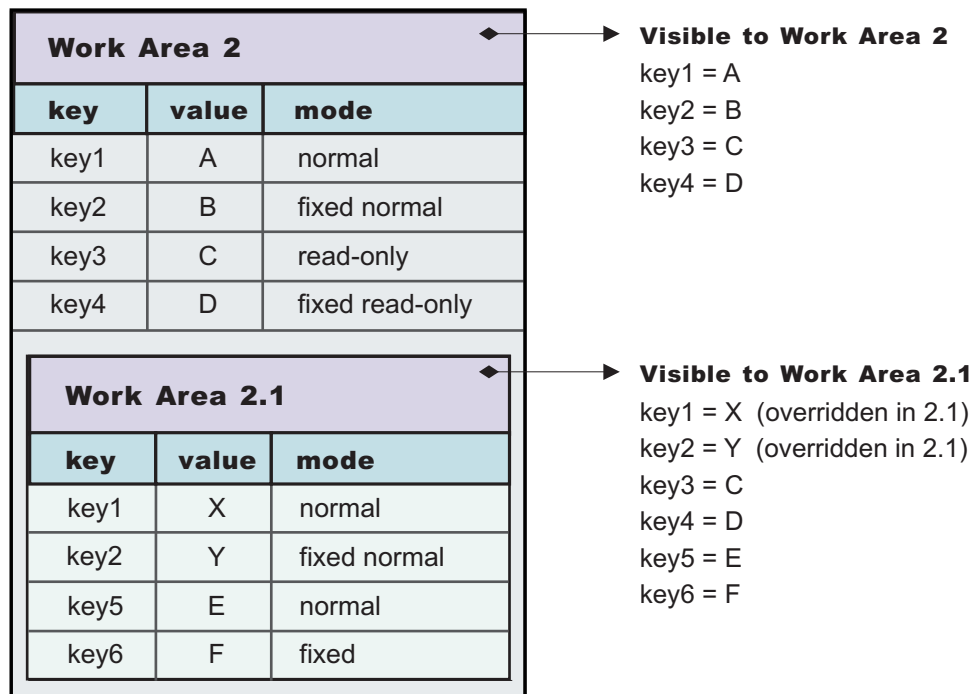


Figure 26. Redefining existing properties in nested work areas

## Distributed work areas

**5.1 +** The propagation of work area context operates differently depending on whether a work area partition is defined as bidirectional or not. In either case all work area context propagates to a target object on a remote invocation. However, whether the context propagates from a target object back to the originator depends on whether a partition is defined as bidirectional.

### Non-bidirectional work area partitions (UserWorkArea partition)

If a remote invocation is issued from a thread associated with a work area, a copy of the work area is automatically propagated to the target object, which can use or ignore the information in the work area as necessary. If the calling application has a nested work area associated with it, a copy of the nested work area and all its ancestors is propagated to the target. The target application can locally modify the information, as allowed by the property modes, by creating additional nested work areas; this information will be propagated to any remote objects it invokes. However, no changes made to a nested work area on a target object are propagated back to the calling object. The caller's work area is unaffected by changes made in the remote method.

## **Bidirectional work area partitions** 5.1+

If a remote invocation is issued from a thread associated with a work area, a copy of the work area is automatically propagated to the target object, which can use or ignore the information in the work area as necessary. If the calling application has a nested work area associated with it, a copy of the nested work area and all its ancestors is propagated to the target. The target application can locally modify the information, as allowed by the property modes, this information is propagated to any remote objects it invokes. In a partition that is not defined as bidirectional, a target application must begin a nested work area before making changes to the imported work area. However, if a partition is defined as bidirectional, a target application need not begin a nested work area before operating on an imported work area. By not beginning a nested work area, any new context set into the work area, or any context changes made by the target application, will not only propagate on future remote invocations but will also propagate back to the originating application (that is, the one who initiated the remote invocation) thus allowing bidirectional propagation of work area context. If the target application does not want new or changed context to propagate back to the originating application, then the target application must begin a nested work area to scope the context to its process. However, the new or changed context in the nested work area will propagate on any future remote invocation the target application may make.

## **WorkArea service: Special considerations**

Developers who use work areas should consider the following issues that could potentially cause problems: interoperability between the EJB and CORBA programming models; and the use of work areas with Java's Abstract Windowing Toolkit.

### **EJB and CORBA interoperability**

Although the WorkArea service can be used across the EJB and CORBA programming models, many composed data types cannot be successfully used across those boundaries. For example, if a SimpleSampleCompany instance is passed from the WebSphere environment into a CORBA environment, the CORBA application can retrieve the SimpleSampleCompany object encapsulated within a CORBA Any object from the work area, but it cannot extract the value from it. Likewise, an IDL-defined struct defined within a CORBA application and set into a work area will not be readable by an application using the UserWorkArea class. Applications can avoid this incompatibility by directly setting only primitive types, like integers and strings, as values in work areas, or by implementing complex values with structures designed to be compatible, like CORBA valuetypes. Also, CORBA Anys that contains either the tk\_null or tk\_void typecode can be set into the work area by using the CORBA interface, but the work-area specification cannot allow the J2EE implementation to return null on a lookup that retrieves these CORBA-set properties without incorrectly implying that there is no value set for the corresponding key. If a J2EE application tries to retrieve CORBA-set properties that are non-serializable, or contain CORBA nulls or void references, the `com.ibm.websphere.workarea.IncompatibleValue` exception is raised.

### **Using work areas with Java's Abstract Windowing Toolkit (AWT)**

Work areas must be used cautiously in applications that use Java's Abstract Windowing Toolkit (AWT). The AWT implementation is multithreaded, and work areas begun on one thread are not available on another. For example, if a program begins a work area in response to an AWT event, such as pressing a button, the work area might not be available to any other part of the application after the execution of the event completes.

## **Work area service performance considerations**

The work area service is designed to address complex data passing patterns that can quickly grow beyond convenient maintenance. A work area is a note pad that is accessible to any client capable of looking up Java Naming Directory Interface (JNDI). Once a work area is established, data can be placed there for future use in any subsequent method calls to both remote and local resources.

You can utilize a work area when a large number of methods require common information or if information is only needed by a method that is significantly further down the call graph. The former avoids the need for complex parameter passing models where the number of arguments passed becomes excessive and hard to maintain. You can improve application function by placing the information in a work area and subsequently accessing it independently in each method, eliminating the need to pass these parameters from method to method. The latter case also avoids unnecessary parameter passing and helps to improve performance by reducing the cost of marshalling and de-marshalling these parameters over the Object Request Broker (ORB) when they are only needed sparsely throughout the call graph.

When attempting to maximize performance by using a work area, cache the `UserWorkArea` retrieved from JNDI wherever it is accessed. You can reduce the amount of time spent looking up information in JNDI by retrieving it once and keeping a reference for the future. JNDI look up takes time and can be costly.

Additional caching mechanisms available to a user defined partition are defined by the configuration property, "Deferred Attribute Serialization". This mechanism will attempt to minimize the number of serialization and deserialization calls. See "Work area partition service" on page 1133 for further explanation of this configuration attribute.

Other configuration parameters that may affect the performance of the work area server are "maxSendSize" and "maxReceiveSize". Setting these two values to 0 (zero) effectively turns off the policing of the size of context that can be sent in a work area. This may enhance performance depending on the number of nested work areas an application uses. In applications that use only one work area, the performance enhancement may be negligible. In applications that have a large number of nested work areas, there may be a performance enhancement. However, a user must note that by turning off this policing it is possible that an extremely large amount of data might be sent to a server.

Performance is degraded if you use a work area as a direct replacement to passing a single parameter over a single method call. The reason is that you incur more overhead than just passing that parameter between method calls. Although the degradation is usually within acceptable tolerances and scales similarly to passing parameters with regard to object size, consider this potential problem before utilizing the service. As with most functional services, intelligent usage of work areas yields the best results.

The work area service is a tool to simplify the job of passing information from resource to resource, and in some cases can improve performance by reducing the overhead associated with parameter passing when the information is only sparsely accessed within the call graph. Caching the instance retrieved from JNDI is important to effectively maximize performance during runtime.

---

## Developing applications that use work areas

Applications interact with the work area service by using the `UserWorkArea` interface and its implementation. This interface defines all of the methods used to create, manipulate, and complete work areas:

1. Access the partition by either:
  - "Accessing the WorkArea service" on page 1124, to access the `UserWorkArea` partition.
  - "Accessing a user defined work area partition" on page 1140, to access a user defined work area.

The following steps use the `UserWorkArea` partition as an example, however a user defined partition can be used in the same way.

2. Beginning a work area.
3. Setting properties in a work area.
4. Using a work area to manage local work.
5. Completing a work area.

An example application, the WorkArea SimpleSample application, is used throughout this documentation to illustrate these tasks

## UserWorkArea interface

Applications interact with the WorkArea service by implementing the UserWorkArea interface. This interface, shown below, defines all of the methods used to create, manipulate, and terminate work areas:

```
package com.ibm.websphere.workarea;

public interface UserWorkArea {
 void begin(String name);
 void complete() throws NoWorkArea, NotOriginator;

 String getName();
 String[] retrieveAllKeys();
 void set(String key, java.io.Serializable value)
 throws NoWorkArea, NotOriginator, PropertyReadOnly;
 void set(String key, java.io.Serializable value, PropertyModeType mode)
 throws NoWorkArea, NotOriginator, PropertyReadOnly;
 java.io.Serializable get(String key);
 PropertyModeType getMode(String key);
 void remove(String key)
 throws NoWorkArea, NotOriginator, PropertyFixed;
}
```

**Note:** EJB applications can use the UserWorkArea interface only within the implementation of methods in the remote interface; likewise, servlets can use the interface only within the service method of the HTTPServlet class. Use of work areas within any lifecycle method of a servlet or enterprise bean is considered a deviation from the work area programming model and is not supported.

### Exceptions

The WorkArea service defines the following exceptions for use with the UserWorkArea interface:

#### NoWorkArea

Thrown when a request requires an associated work area but none is present.

#### NotOriginator

Raised when a request attempts to manipulate the contents of an imported work area.

#### PropertyReadOnly

Raised when a request attempts to modify a read-only or fixed read-only property.

#### PropertyFixed

Raised by the remove method when the designated property has one of the fixed modes.

## Example: WorkArea SimpleSample application

In this example, the client creates a work area and inserts two properties into the work area: a site identifier and a priority. The site-identifier is set as a read-only property; the client does not allow recipients of the work area to override the site identifier. This property consists of the key company and a static instance of a SimpleSampleCompany object. The priority property consists of the key priority and a static instance of a SimpleSamplePriority object. The object types are defined as shown in the following code example

```
public static final class SimpleSampleCompany {
 public static final SimpleSampleCompany Main;
 public static final SimpleSampleCompany NewYork_Sales;
 public static final SimpleSampleCompany NewYork_Development;
 public static final SimpleSampleCompany London_Sales;
 public static final SimpleSampleCompany London_Development;
}

public static final class SimpleSamplePriority {
 public static final SimpleSamplePriority Platinum;
 public static final SimpleSamplePriority Gold;
}
```

```

 public static final SimpleSamplePriority Silver;
 public static final SimpleSamplePriority Bronze;
 public static final SimpleSamplePriority Tin;
}

```

The client then makes an invocation on a remote object. The work area is automatically propagated; none of the methods on the remote object take a work area argument. On the remote side, the request is first handled by the SimpleSampleBean; the bean first reads the site identifier and priority properties from the work area. The bean then intentionally attempts, and fails, both to write directly into the imported work area and to override the read-only site-identifier property.

The SimpleSampleBean successfully begins a nested work area, in which it overrides the client's priority, then calls another bean, the SimpleSampleBackendBean. The SimpleSampleBackendBean reads the properties from the work area, which contains the site identifier set in the client and priority set in the SimpleSampleBean. Finally, the SimpleSampleBean completes its nested work area, writes out a message based on the site-identifier property, and returns.

The implementation of this application is discussed in the topic, [Developing applications that use work areas](#).

## Accessing the WorkArea service

The WorkArea service provides a JNDI binding to an implementation of the UserWorkArea interface under the name `java:comp/websphere/UserWorkArea`. Applications that need to access the service can perform a lookup on that JNDI name, as shown in the following code example:

```

import com.ibm.websphere.workarea.*;
import javax.naming.*;

public class SimpleSampleServlet {
 ...

 InitialContext jndi = null;
 UserWorkArea userWorkArea = null;
 try {
 jndi = new InitialContext();
 userWorkArea = (UserWorkArea)jndi.lookup(
 "java:comp/websphere/UserWorkArea");
 }
 catch (NamingException e) { ... }
}

```

The next step is to use the `begin()` method to create a new work area and associate it with the calling thread, as described in the topic [Beginning a new work area](#).

## Beginning a new work area

Be sure that your client has a reference to the UserWorkArea interface, as described in the topic [Accessing the WorkArea Service](#) or a reference to a user defined partition as defined in “[Accessing a user defined work area partition](#)” on page 1140. The following steps use the UserWorkArea partition as an illustration. However a user defined partition can be used in the exact same way.

Use the `begin()` method to create a new work area and associate it with the calling thread. The `begin()` method takes a string as an argument; the string is used to name the work area. The argument must not be null, which causes the `java.lang.NullPointerException` to be raised. In the following code example, the application begins a new work area with the name SimpleSampleServlet:

```

public class SimpleSampleServlet {
 ...
 try {

```

```

 ...
 userWorkArea = (UserWorkArea)jndi.lookup(
 "java:comp/websphere/UserWorkArea");
}
...

userWorkArea.begin("SimpleSampleServlet");
...
}

```

The `begin()` method is also used to create nested work areas; if a work area is associated with a thread when the `begin()` method is called, the method creates a new work area nested within the existing work area.

The `WorkArea` service makes no use of the names associated with work areas; You can name work areas in any way that you choose. Names are not required to be unique, but the usefulness of the names for debugging is enhanced if the names are distinct and meaningful within the application. Applications can use the `getName()` method to return the name associated with a work area by the `begin()` method.

Using a work area

## Setting properties in a work area

An application with a current work area can insert properties into the work area and retrieve the properties from the work area. The `UserWorkArea` interface provides two `set()` methods for setting properties and a `get()` method for retrieving properties. The two-argument `set()` method inserts the property with the property mode of normal. The three-argument `set()` method takes a property mode as the third argument. (See "Setting property modes", later in this topic.)

Both `set()` methods take the key and the value as arguments. The key is a `String`; the value is an object of the type `java.io.Serializable`. None of the arguments can be null, which causes the `java.lang.NullPointerException` to be raised.

The `WorkArea` `SimpleSample` application uses objects of two classes, the `SimpleSampleCompany` class and the `SimpleSampleProperty` class, as values for properties. The `SimpleSampleCompany` class is used for the site identifier, and the `SimpleSamplePriority` class is used for the priority. These classes are shown in following code example:

```

public class SimpleSampleServlet {
 ...
 userWorkArea.begin("SimpleSampleServlet");

 try {
 // Set the site-identifier (default is Main).
 userWorkArea.set("company",
 SimpleSampleCompany.Main, PropertyModeType.read_only);

 // Set the priority.
 userWorkArea.set("priority", SimpleSamplePriority.Silver);
 }

 catch (PropertyReadOnly e) {
 // The company was previously set with the read-only or
 // fixed read-only mode.
 ...
 }

 catch (NotOriginator e) {
 // The work area originated in another process,
 // so it can't be modified here.
 ...
 }
}

```



```

catch (NoWorkArea e) {
 // There is no work area begun on this thread.
 ...
}

// Do application work.
...
}

```

The `get()` method takes the key as an argument and returns a Java Serializable object as the value associated with the key. For example, to retrieve the value of the company key from the work area, the code example above uses the `get()` method on the work area to retrieve the value.

**Setting property modes.** The two-argument `set()` method on the `UserWorkArea` interface takes a key and a value as arguments and inserts the property with the default property mode of normal. To set a property with a different mode, applications must use the three-argument `set()` method, which takes a property mode as the third argument. The values used to request the property modes are as follows:

- **Normal:** `PropertyModeType.normal`
- **Fixed normal:** `PropertyModeType.fixed_normal`
- **Read-only:** `PropertyModeType.read_only`
- **Fixed read-only:** `PropertyModeType.fixed_readonly`

## Using a work area to manage local work

Be sure that your client has a reference to the `UserWorkArea` interface, as described in the topic [Accessing the WorkArea Service](#) or a reference to a user defined partition as defined in “Accessing a user defined work area partition” on page 1140. The following steps use the `UserWorkArea` partition as an illustration. However a user defined partition can be used in the exact same way.

In a business application that uses work areas, server objects typically retrieve the work area properties and use them to guide local work.

1. Retrieving the name of the active work area This step determines whether the calling thread is associated with a work area.
2. Overriding work area properties. Server objects can override client work area properties by creating their own, nested work area.
3. Retrieving properties from a work area
4. Retrieving a list of all keys in a work area
5. Querying the mode of a work area property
6. Deleting a work area property
7. Completing a work area

The server side of the `WorkArea SimpleSample` application accepts remote invocations from clients. With each remote call, the server also gets a work area from the client if the client has created one. The work area is propagated transparently. None of the remote methods includes the work area on its argument list.

In the example application, the server objects use the work area interface for demonstration purposes only. For example, the `SimpleSampleBean` intentionally attempts to write directly to an imported work area, which triggers the `NotOriginator` exception. Likewise, the bean intentionally attempts to mask the read only `SimpleSampleCompany`, which triggers the `PropertyReadOnly` exception. The `SimpleSampleBean` also nests a work area and successfully overrides the priority property before invoking the `SimpleSampleBackendBean`. A true business application would extract the work area properties and use them to guide the local work. The `SimpleSampleBean` mimics this by writing a message that function is denied when a request emanates from a sales environment.

## Retrieving the name of the active work area

Applications use the `getName()` method on the `UserWorkArea` interface to retrieve the name of the current work area. This is the recommended method for determining whether the thread is associated with a work area; if the thread is not associated with a work area, the `getName()` method returns null. In the following code example, the name of the work area corresponds to the name of the class in which the work area was begun.

```
public class SimpleSampleBeanImpl implements SessionBean {

 ...

 public String [] test() {
 // Get the work-area reference from JNDI.
 ...

 // Retrieve the name of the work area. In this example,
 // the name is used to identify the class in which the
 // work area was begun.
 String invoker = userWorkArea.getName();
 ...
 }
}
```

## Overriding work area properties

Work areas are inherently associated with the process that creates them. In the sample application, the client begins a work area and sets into it the site-identifier and priority properties. This work area is propagated to the server when the client makes a remote invocation.

Applications nest work areas in order to temporarily override properties imported from a client process. The nesting mechanism is automatic; invoking `begin` on the `UserWorkArea` interface from within the scope of an existing work area creates a nested work area that inherits the properties from the enclosing work area. Properties set into the nested work area are strictly associated with the process in which the work area was begun; the nested work area must be completed within the process that created them. If a work area is not completed by the creating process, the work-area facility terminates the work area when the process exits. After a nested work area is completed, the original view of the enclosing work area is restored. However, the view of the complete set of work areas associated with a thread cannot be decomposed by downstream processes.

Applications set properties into a work area using property modes to ensure that a particular property is fixed (not removable) or read-only (not overrideable) within the scope of the given work area.

In the following code example, the server-side sample bean attempts to write directly to the imported work area; because the `UserWorkArea` partition is not defined to be bidirectional, this action is not permitted, and the `NotOriginator` exception is thrown. When the `UserWorkArea` partition is not defined as bidirectional, the sample bean must begin its own work area in order to override any imported properties, as shown in the second code example. If a work area in a user defined partition is used and is defined as bidirectional, this bean can set context into the work area before beginning another work area. This context set in the bidirectional case will propagate back to the caller. See "Bidirectional propagation" on page 1139 for additional information.

```
public class SimpleSampleBeanImpl implements SessionBean {

 public String [] test() {
 ...
 String invoker = userWorkArea.getName();

 try {
 userWorkArea.set("key", "value");
 }
 catch (NotOriginator e) {
 ...
 }
 }
}
```

```

 }
 ...
}
}

```

The following code example demonstrates beginning a nested work area, using the name of the creating class to identify the nested work area.

```

public class SimpleSampleBeanImpl implements SessionBean {

 public String [] test() {
 ...
 String invoker = userWorkArea.getName();
 try {
 userWorkArea.set("key", "value");
 }
 catch (NotOriginator e) {
 }

 // Begin a nested work area. By using the name of the creating
 // class as the name of the work area, we can avoid having
 // to explicitly set the name of the creating class in
 // the work area.
 userWorkArea.begin("SimpleSampleBean");

 ...
 }
}

```

In the example application, the client sets the site-identifier property as read-only; that guarantees that the request will always be associated with the client's company identity. A server cannot override that value in a nested work area. In the following code example, the SimpleSampleBean attempts to change the value of the site-identifier property in the nested work area it created.

```

public class SimpleSampleBeanImpl implements SessionBean {

 public String [] test() {
 ...

 String invoker = userWorkArea.getName();
 try {
 userWorkArea.set("key", "value");
 }
 catch (NotOriginator e) {
 }

 // Begin a nested work area.
 userWorkArea.begin("SimpleSampleBean");

 try {
 userWorkArea.set("company",
 SimpleSampleCompany.London_Development);
 }
 catch (NotOriginator e) {
 }
 ...
 }
}

```

## Retrieving work area properties

Properties can be retrieved from a work area by using the `get()` method. This method is intentionally lightweight; there are no declared exceptions to handle. If there is no active work area, or if there is no such property set in the current work area, the `get()` method returns null.

**Note:** The `get()` method can raise a `NotSerializableError` in the relatively rare scenario in which CORBA clients set composed data types and invoke enterprise-bean interfaces.

The following example shows the retrieval of the site-identifier and priority properties by the `SimpleSampleBean`. Recall that one property was set into an outer work area by the client, and the other property was set into the nested work area by the server-side bean; the nesting is transparent to the retrieval of the properties.

```
public class SimpleSampleBeanImpl implements SessionBean {

 public String [] test() {
 ...

 // Begin a nested work area.
 userWorkArea.begin("SimpleSampleBean");
 try {
 userWorkArea.set("company",
 SimpleSampleCompany.London_Development);
 }
 catch (NotOriginator e) {
 }

 SimpleSampleCompany company =
 (SimpleSampleCompany) userWorkArea.get("company");
 SimpleSamplePriority priority =
 (SimpleSamplePriority) userWorkArea.get("priority");
 ...
 }
}
```

## Retrieving a list of all keys in a work area

The `UserWorkArea` interface provides the `retrieveAllKeys()` method for retrieving a list of all the keys visible from a work area. This method takes no arguments and returns an array of strings. This method returns null if there is no work area associated with the thread. If there is an associated work area containing no properties, the method returns an array of size 0.

## Querying the mode of a work area property

The `UserWorkArea` interface provides the `getMode()` method for determining the mode of a specific property. This method takes the property's key as an argument and returns the mode as a `PropertyModeType` object. (See *Setting property modes* for more information on names of mode types.) If the specified key does not exist in the work area, the method returns `PropertyModeType.normal`, indicating that the property can be set and removed without error.

## Deleting a work area property

The `UserWorkArea` interface provides the `remove()` method for deleting a property from the current scope of a work area. If the property was initially set in the current scope, removing it deletes the property. If the property was initially set in an enclosing work area, removing it deletes the property until the current scope is completed. When the current work area is completed, the deleted property is restored.

The `remove()` method takes the property's key as an argument. Only properties with the modes `normal` and `read-only` can be removed. Attempting to remove a fixed property causes the `PropertyFixed` exception to be thrown. Attempting to remove properties in work areas that originated in other processes causes the `NotOriginator` exception to be thrown.

## Completing a work area

After an application has finished using a work area, it must complete the work area by calling the `complete()` method on the `UserWorkArea` interface. This terminates the association with the calling thread

and destroys the work area. If the complete method is called on a nested work area, the nested work area is terminated and the parent work area becomes the current work area. If there is no work area associated with the calling thread, a NoWorkArea exception is thrown.

Every work area must be completed, and work areas can be completed only by the originating process. For example, if a server attempts to call the complete() method on a work area that originated in a client, a NotOriginator exception is thrown. Work areas created in a server process are never propagated back to an invoking client process.

**Note:** The WorkArea service claims full local-remote transparency. Even if two beans happen to be deployed in the same server, and therefore the same JVM and process, a work area begun on an invocation from another is completed and the bean in which the request originated is always in the same state after any remote call.

The following code example shows the completion of the work area created in the client application.

```
public class SimpleSampleServlet {
 ...
 userWorkArea.begin("SimpleSampleServlet");
 userWorkArea.set("company",
 SimpleSampleCompany.Main, PropertyModeType.read_only);
 userWorkArea.set("priority", SimpleSamplePriority.Silver);
 ...

 // Do application work.
 ...

 // Terminate the work area.
 try {
 userWorkArea.complete();
 }

 catch (NoWorkArea e) {
 // There is no work area associated with this thread.
 ...
 }

 catch (NotOriginator e) {
 // The work area was imported into this process.
 ...
 }
 ...
}
```

The following code example shows the sample application completing the nested work area it created earlier in the remote invocation.

```
public class SimpleSampleBeanImpl implements SessionBean {

 public String [] test() {
 ...

 // Begin a nested work area.
 userWorkArea.begin("SimpleSampleBean");
 try {
 userWorkArea.set("company",
 SimpleSampleCompany.London_Development);
 }
 catch (NotOriginator e) {
 }

 SimpleSampleCompany company =
 (SimpleSampleCompany) userWorkArea.get("company");
 SimpleSamplePriority priority =
 (SimpleSamplePriority) userWorkArea.get("priority");
 }
}
```

```

 // Complete all nested work areas before returning.
 try {
 userWorkArea.complete();
 }
 catch (NoWorkArea e) {
 }
 catch (NotOriginator e) {
 }
}
}

```

---

## Managing the work area service

The WorkArea service is managed using the administrative console. There are two administrative tasks associated with work areas:

- Enabling the WorkArea service. The WorkArea Service is enabled by default on both clients and servers.
- Managing the size of work areas. Applications can set maximum sizes on each work area to be sent and to be accepted.

### Enabling the WorkArea service

For an application to take advantage of work areas, the WorkArea service must be enabled for both clients and servers. In both cases, the service is enabled by default.

1. Enable (or disable) the use of work areas on a server:
  - a. Start the administrative console.
  - b. Select **Servers > server\_name > WorkArea Service**.
  - c. Select or clear the **Startup** checkbox. This specifies whether or not the server should automatically start the WorkArea service when the server starts.
2. Enable (or disable) the use of work areas on a client: Set the `com.ibm.websphere.workarea.enabled` property to `TRUE` or `FALSE` before starting the client. For example, edit the `launchClient` script in the `$WAS_HOME/bin` directory and add the following to the Java invocation:

```
-Dcom.ibm.websphere.workarea.enabled=false
```

### WorkArea service settings

Use this page to manage the work area service.

The work area service manages the scope and implicit propagation of application context.

To view this administrative console page, click **Servers > Application Servers > server\_name > Work Area Service** .

#### **Startup:**

Specifies whether the server will attempt to start the work area service.

**Default**

Selected

**Range****Selected**

When the application server starts, it attempts to start the work area service automatically.

**Cleared**

The server does not try to start the work area service. If work areas are to be used on this application server, the system administrator must start the service manually or select this property then restart the server.

**Maximum Send Size:**

Specifies the maximum size of data that can be sent within a single work area.

<b>Data type</b>	Integer
<b>Units</b>	Bytes
<b>Default</b>	32767
<b>Range</b>	-1 to no limit
	<b>-1</b> Default.
	<b>0</b> No limit.

**Maximum Receive Size:**

Specifies the maximum size of data that can be received by a single work area.

<b>Data type</b>	Integer
<b>Units</b>	Bytes
<b>Default</b>	32767
<b>Range</b>	-1 to no limit
	<b>-1</b> Default.
	<b>0</b> No limit.

## Managing the size of work areas

Applications can set maximum sizes on each work area to be sent or received. By default, the maximum size of a work area that is sent by a client and received, then possibly re-sent, by a server is 32,768 bytes. You can change this size as described in this topic.

1. Change the size of the work area that can be sent or received by a server:
  - a. Start the administrative console.
  - b. Select **Servers > server\_name > WorkArea Service**.
  - c. Enter a new value in the **maxSendSize** field to modify the size of the work area that this server can send, or enter a new value in the **maxReceiveSize** field to modify the size of the work area that this server can accept.
2. Change the size of the work area that can be sent by a client: Set the `com.ibm.websphere.workarea.maxSendSize` property to the desired number of bytes before starting the client. This can be done in several ways. For example, to set the maximum size to 10,000 bytes, edit the `launchClient` script in the `$WAS_HOME/bin` directory and add the following to the Java invocation:

```
-Dcom.ibm.websphere.workarea.maxSendSize=10000
```

The maximum size that can be specified is determined by the maximum value expressible in the Java Integer data type, 2,147,483,647. The smallest maximum size that can be specified is 1. Using a maximum size of 1 byte effectively means that no requests associated with the work area can leave the



system or enter another system. A value of 0 means that no limit is imposed. A value of -1 means that the default value is to be honored. The default value is also used if an invalid value or a malformed property is specified.

---

## Configuring work area partitions on the server

The work area partition service extends the work area service by allowing the creation of multiple work areas with more configuration options. Follow these steps to create and configure a work area partition:

1. Start the administrative console.
2. Click **Servers > Application Servers > *server\_name* > Work Area Partition Service**.
3. Click **New**.
4. On the settings page for work area partitions, specify values such as the partition name, `maxSendSize` and `maxReceiveSize`, then click OK.
5. Save the new configuration and restart the server to apply the new configuration

You have created a work area partition

Retrieve the partition through the work area partition manager interface and use it as defined by the work area service and the work area service interface. See the topic, “Example: Work area partition manager” on page 1136, for an example.

## Work area partition service

The work area partition service is an extension of the work area service that allows the creation of multiple custom work areas. The work area partition service is an optional service to users. Any user that currently uses the work area service and the `UserWorkArea` partition can continue using it in the same manner. The `UserWorkArea` partition is created automatically (if it has not been disabled) by the work area partition service. By allowing a user the option to create their own work area partition through the work area partition service, they can have more control over configuration and access to their partition.

Unlike the `UserWorkArea` partition, which is publicly known, work areas created by the work area partition service are accessible to, and known only by the creator. However, the work area partition service does not strictly enforce that a partition is accessed and/or operated on exclusively by the partition creator. There are no limitations should the creator want to publish their work area partition and make it publicly available by binding their partition reference in java naming or by other means. However, the work area partition service does try to hide a partition as much as possible should a user not want others to know about a certain partition. The work area partition service does not allow a person to determine, or query the names of all the partitions that have been created; however, it does not restrict the partitions from being accessed by users other than the creator of that partition. The context of a partition, such as the `UserWorkArea` partition or a user defined partition, is scoped to a single thread and is not accessible by multiple threads.

The work area partition reference returned to a user implements `javax.naming.Referenceable` (as well as `com.ibm.websphere.UserWorkArea`), therefore a user can bind their partition into naming should they want to make their partition publicly available. An alternative to using java naming to bind and access the partition is to use the work area partition manager interface. Anyone can access the work area partition manager interface; therefore, if a user wants to make their partition publicly available, they simply need to publish their partition name. Other users can then call the `getWorkAreaPartition(..)` method on the work area partition manager interface with the published name.

The `WorkAreaPartitionManager.createWorkAreaPartition(..)` method can only be used from a J2EE client. To create a work area partition on the server side, one must use the administrative console. On the server side a work area partition must be created during server startup because each partition needs to be register with the appropriate web and EJB collaborators before the server has started. Custom work area partitions are created by the work area partition service and defined by the `UserWorkArea` interface.

The work area partition service also allows a user to configure partitions with additional properties that are not available on the UserWorkArea partition, such as bidirectional propagation of work area partition context and deferred attribute serialization. These properties are available as configuration properties when creating a partition. The properties are defined as follows:

### **Bidirectional propagation of WorkArea context**

If a remote invocation is issued from a thread associated with a work area, a copy of the work area is automatically propagated to the target object, which can use or ignore the information in the work area as necessary. If the calling application has a nested work area associated with it, a copy of the nested work area and all its ancestors are propagated to the target application. The target application can locally modify the information, as allowed by the property modes, by creating additional nested work areas; this information will be propagated to any remote objects that it invokes.

Whether context changes will propagate back to a calling application from a remote application depends on the configuration of the work area partition. If a user creates a partition to be bidirectional (selects the Bidirectional property during creation), changes made by a remote application will propagate back to the calling application, meaning that changes made to the work area context by a downstream process will propagate back up stream. The UserWorkArea partition is not configured (and can never be configured) to be bidirectional; therefore context changes will only flow to downstream processes and will not propagate back upstream. See “Bidirectional propagation” on page 1139 for further explanation.

### **Deferred attribute serialization of WorkArea context**

By default, on each set operation the attribute set into a work area is automatically serialized by the work area service. On each subsequent get operation on that same attribute it is deserialized and returned to the requester. This gives the work area service complete control of the attribute such that any changes to a mutable object are not reflected in the work area’s copy of the attribute unless a user specifically resets the attribute into the work area. However, this can potentially lead to excessive serialization and deserialization. Excessive serialization and deserialization can result in observable performance degradation under heavy load. The deferred attribute serialization configuration property is a caching feature that reduces serialization and deserialization operations. When deferred attribute serialization is enabled (by selecting the Deferred Attribute Serialization field during creation) in a client or server process, attributes set into the work area Service are not automatically serialized during the set operation. Rather, a reference to the attribute is stored in the work area. If the attribute is mutable, then changes to the object will be reflected in the work area’s reference to that attribute. When a get operation is performed on that attribute, the reference to that object is returned and no deserialization is performed. Attributes are not actually serialized until the thread with which the attribute is associated makes a remote IIOp invocation. At that point the attribute is serialized and the serialized form of the attribute is cached. If the attribute is not reset into the work area, changes to the original attribute are still reflected within the attribute contained within the work area because the work area still holds a cached reference to the original object. However, if the work area has not been told that the attribute has changed by resetting the attribute into the work area, subsequent remote requests continues to use the cached serialized version of the attribute and direct changes to the mutable attribute are not propagated. This is an important distinction between enabling and not enabling the deferred attribute serialization configuration property and a user must pay close attention to this difference and how mutable objects are handled when enabling deferred attribute serialization. The work area service releases cached references and cached serialized versions of attributes when any of the following occur:

- An attribute is reset or removed.
- The work area is explicitly completed by the application.
- Server component ends execution of the request during which the work area was begun.
- Client process which began the work area terminates.

## Partition context propagation across process boundaries

Work area context automatically propagates from client to server when a client makes a remote call to a server. If a client is configured with, for example, three different work area partitions when it makes a remote call to a server, server1; the context associated with each partition on the client thread propagates to server1. If the same three partitions reside (have been created) on server1, the context is demarshaled to the appropriate partition. However, if none or only a few of the three partitions have been created on server1, only the context associated with a partition that is resident on both the client and server is demarshaled. The context associated with a partition that is not resident on server1 is still resident on server1 but will not be accessible. The context associated with partitions that are not resident on server1 must remain resident on server1 in case another remote call is made to a different server. Going one step further, if server1 makes a call to yet another server, server2 and assume server2 has created all the same partitions that the client has, server2 receives the context for the partitions that were not resident on server1. Any partitions that reside on server1 that did not reside on the client, now have its context propagated to server2.

## WorkAreaPartitionManager interface

Applications interact with the work area partition service by using the WorkAreaPartionManager interface. A user can retrieve an instance of the WorkAreaPartitionManager out of naming and use the methods defined below. An implementation of the WorkAreaPartitionManager is bound in java naming at "java:comp/websphere/WorkAreaPartitionManager". This interface is responsible for creating, retrieving, and manipulating work area partitions:

```
package com.ibm.websphere.workarea;

import com.ibm.websphere.workarea.UserWorkArea;
import com.ibm.websphere.workarea.PartitionAlreadyExistsException;
import com.ibm.websphere.workarea.NoSuchPartitionException;
import java.util.Properties;

public interface WorkAreaPartitionManager {

 public UserWorkArea getWorkAreaPartition(String partitionName) throws NoSuchPartitionException;

 public UserWorkArea createWorkAreaPartition(String partitionName, Properties props) throws
 PartitionAlreadyExistsException, java.lang.IllegalAccessException;
}
```

Programmatically creating a work area partition (through the createWorkAreaPartition(..) method) is only available on the J2EE client. To create a work area partition on the server, one must use the WebSphere Administrative Console. All partitions in a server process must be created before server start up is complete to allow the work area service to register with the appropriate container collaborators. Therefore, calling createWorkAreaPartition(..) in a server process after the server has started results in a java.lang.IllegalAccessException. The createWorkAreaPartition(..) method can be called in a J2EE process at any time.

### Exceptions

The work area partition service defines the following exceptions for use with the WorkAreaPartitionManager interface:

#### **PartitionAlreadyExistsException**

This exception is raised by the createWorkAreaPartition(..) method on the WorkAreaPartitionManager implementation if a user tries to create a work area partition with a partition name that already exists. Partition names must be unique.

#### **NoSuchPartitionException**

This exception is raised by the getWorkAreaPartition(..) method on the WorkAreaPartitionManager implementation if a user requests a work area partition with a partition name that does not exist.

## java.lang.IllegalAccessException

This exception is raised by the createWorkAreaPartition(..) method on the WorkAreaPartitionManager implementation if a user tries to create a work area partition during runtime on a server process. This method can only be used on a J2EE client process. On the server process a partition must be created using the Administrative Console.

## Example: Work area partition manager

The example below demonstrates the use of the work area partition manager interface. The sample illustrates how to create and retrieve a work area partition programmatically. Please note that programmatically creating a work area partition is only available on the Java 2 platform, Enterprise Edition (J2EE) client. To create a work area partition on the server one must use the WebSphere administrative console. See "Work area partition service" on page 1133 for configuration parameters available to configure a partition.

```
import com.ibm.websphere.workarea.WorkAreaPartitionManager;
import com.ibm.websphere.workarea.UserWorkArea;
import com.ibm.websphere.workarea.PartitionAlreadyExistsException;
import com.ibm.websphere.workarea.NoSuchPartitionException;
import java.lang.IllegalAccessError;
import java.util.Properties;
import javax.naming.InitialContext;

//This sample demonstrates how to retrieve an instance of the
//WorkAreaPartitionManager implementation and how to use that
//instance to create a WorkArea partition and retrieve a partition.
//NOTE: Creating a partition in the way listed below is only available
//on a J2EE client. To create a partition on the server use the
//WebSphere administrative console. Retrieving a WorkArea
//partition is performed in the same way on both client and server.

public class Example {

 //The name of the partition to create/retrieve
 String partitionName = "myPartitionName";
 //The name in java naming the WorkAreaPartitionManager instance is bound to
 String jndiName = "java:comp/websphere/WorkAreaPartitionManager";

 //On a J2EE client a user would create a partition as follows:
 public UserWorkArea myCreate(){
 //Variable to hold our WorkAreaPartitionManager reference
 WorkAreaPartitionManager partitionManager = null;
 //Get an instance of the WorkAreaPartitionManager implementation
 try {
 InitialContext initialContext = new InitialContext();
 partitionManager = (WorkAreaPartitionManager) initialContext.lookup(jndiName);
 } catch (Exception e) { }

 //Set the properties to configure our WorkArea partition
 Properties props = new Properties();
 props.put("maxSendSize","12345");
 props.put("maxReceiveSize","54321");
 props.put("Bidirectional","true");
 props.put("DeferredAttributeSerialization","true");

 //Variable used to hold the newly created WorkArea Partition
 UserWorkArea myPartition = null;

 try{
 //This is the way to create a partition on the J2EE client. Use the
 //WebSphere Administrative Console to create a WorkArea Partition
 //on the server.
 myPartition = partitionManager.createWorkAreaPartition(partitionName,props);
 }
 catch (PartitionAlreadyExistsException e){ }
 }
}
```

```

 catch (IllegalAccessException e){ }

 return myPartition;
 }

 // . . .

 //In order to retrieve a WorkArea partition at some time later or
 //from some other class, do the following (from client or server):
 public UserWorkArea myGet(){
 //Variable to hold our WorkAreaPartitionManager reference
 WorkAreaPartitionManager partitionManager = null;
 //Get an instance of the WorkAreaPartitionManager implementation
 try {
 InitialContext initialContext = new InitialContext();
 partitionManager = (WorkAreaPartitionManager) initialContext.lookup(jndiName);
 } catch (Exception e) { }

 //Variable used to hold the retrieved WorkArea partition
 UserWorkArea myPartition = null;
 try{
 myPartition = partitionManager.getWorkAreaPartition(partitionName);
 }catch(NoSuchPartitionException e){ }

 return myPartition;
 }
}

```

## Work area partition collection

Use this page to manage the work area service.

The work area partition service supports the definition of custom work area partitions.

To view this administrative console page, click **Servers > Application Servers > *server\_name* > Work Area Partition Service > Work Area Partitions**.

### Name

Specifies the name of the work area partition. This name must be unique and is used to retrieve the partition.

### Description

Specifies the description of the work area partition.

### Startup

Specifies whether the server attempts to start the specified service when the server starts.

### Bidirectional

Permits applications to modify a work area's context that is imported by a J2EE request; modified properties are propagated back to the requestor's environment. This option is disabled by default.

### Maximum send size

The maximum size of data that can be sent within a single work area. (0 = no limit; -1 = default)

### Maximum receive size

The maximum size of data that can be received within a single work area. (0 = no limit; -1 = default)

### Deferred attribute serialization

Specifies whether attribute serialization is deferred until the work area is propagated on a remote invocation.

## Work area partition settings

Use this page to modify the work area service settings.

The work area partition service supports the definition of custom work area partitions.

To view this administrative console page, click **Servers > Application Servers > server\_name > Work Area Partition Service > Work Area Partitions > Work\_area\_partion\_name**.

### **Name:**

Specifies the name of the work area partition. This unique name is used to retrieve the partition.

Extended information about the field, if needed, can be added here.

**Data type** String

### **Description:**

Specifies the description of the work area partition.

**Data type** String

### **Startup:**

Specifies whether the server attempts to start the specified service when the server starts.

**Data type** Check box  
**Default** Selected

### **Bidirectional:**

Permits applications to modify a work area's context that is imported by a J2EE request; modified properties are propagated back to the requestor's environment. This option is disabled by default.

**Data type** Check box  
**Default** Cleared

### **Maximum send size:**

The maximum size of data that can be sent within a single work area. (0 = no limit; -1 = default)

**Data type** Integer  
**Units** Milliseconds|Seconds|Hours|Etc.  
**Default** 32768  
**Range** 0 (no limit) to -1 (default)

### **Maximum receive size:**

The maximum size of data that can be received within a single work area. (0 = no limit; -1 = default)

**Data type** Integer  
**Units** Milliseconds|Seconds|Hours|Etc.  
**Default** 32768

Range

0 (no limit) to -1 (default)

**Deferred attribute serialization:**

Specifies whether attribute serialization is deferred until the work area is propagated on a remote invocation. This option is disabled by default.

Data type

Check box

Default

Cleared

## Bidirectional propagation

### Example: Bidirectional propagation of work area context

Whether context changes propagate back to a calling application from a remote application depends on the configuration of the work area partition. If a user creates a bidirectional partition, changes made by a remote application propagate back to the calling application. In other words, changes made to the work area context by a downstream process propagate back up stream. Figure 1 illustrates distribution of work area context on a remote call when the partition containing the given work area is configured for bidirectional propagation of its work area context. For this illustration, the client and server must have created a partition with the same name.

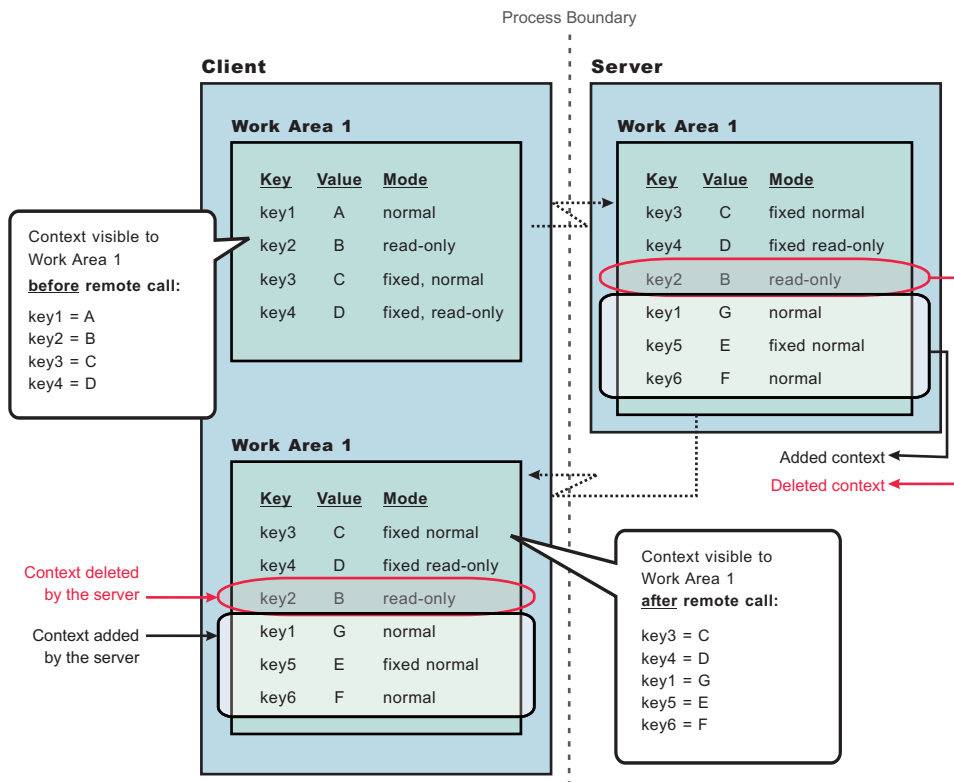


Figure 27. Figure 1

As Figure 1 shows, when the client makes a remote call to the server, the server receives the context set by the client process. The server then can make changes to this context or add to it. In this illustration, the server overwrites the value at **key1**, removes the property at **key2**, and adds two new properties at **key5** and **key6**. When the server application returns to the client, the work area context is propagated back to



the client and demarshalled. The current work area is then updated with the new context. Note, that if the partition is not configured as bidirectional, and the server tries to change or remove context in work area, "Work Area 1", it will receive a `com.ibm.websphere.workarea.NotOriginator` exception since the client was the originator of the work area. The server can retrieve the context in "Work Area 1". This is the main distinction between bidirectional propagation of context and non-bidirectional propagation.

### Bidirectional propagation of nested work area context

If a remote application needs to add context to a work area that is only used by itself or any other remote objects, the remote application should begin another work area. By beginning a new work area, the new context added is scoped to that application and does not flow back to the calling application. The major benefit of nesting work areas is that nesting work areas allows an application to scope work area context to a given application. Taking the above illustration one step further, if the server has begun a work area before overwriting the value at **key1**, removing the property at **key2**, or adding new properties at **key5** and **key6**; those changes would not have propagated back to the client. This is shown in Figure 2. You can also see from this figure that the client does not receive the context from the nested work area started by the server.

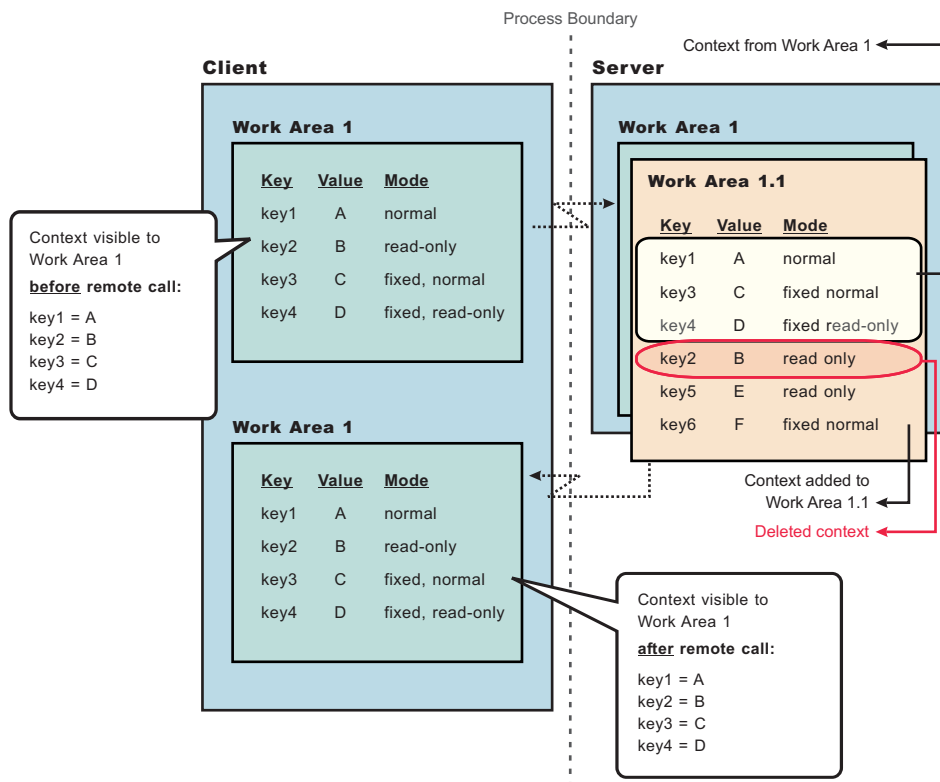


Figure 28. Figure 2

### Accessing a user defined work area partition

The work area partition service provides a Java Naming and Directory Interface (JNDI) binding to an implementation of the work area partition manager interface under the name `java:comp/websphere/WorkAreaPartitionManager`. Applications that need to access their partition can perform a lookup on that JNDI name and then use the `getWorkAreaPartition()` method on the work area partition manager, as shown in the following code example:

```

import com.ibm.websphere.workarea.*;
import javax.naming.*;

public class SimpleSampleServlet {
 ...

 //Variable to hold our WorkAreaPartitionManager implementation
 WorkAreaPartitionManager partitionManager = null;
 try {
 InitialContext initialContext = new InitialContext();
 partitionManager = (WorkAreaPartitionManager)
 initialContext.lookup("java:comp/websphere/UserWorkArea");
 } catch (Exception e) {...}

 //Variable used to hold the retrieved WorkArea Partition
 UserWorkArea myPartition = null;
 try{
 myPartition = partitionManager.getWorkAreaPartition(partitionName);
 }catch(NoSuchPartitionException e){...}
}

```

The next step is to use the `begin()` method to create a new work area and associate it with the calling thread, as described in the topic [Beginning a new work area](#).



---

## Chapter 23. Working with the Event Programming Model in WebSphere

How to use the Common Event Infrastructure in the WebSphere programming environment.

The Event Programming Model is part of the wider programming model in WebSphere. The Event Programming Model provides a standard mechanism for the actions of applications to be recorded as events, or to trigger work in other applications. Events can also be generated at significant points from within processes, providing a powerful tool for analyzing the execution of the process. The Event Programming Model provides the following facilities :

- Standard interfaces and services for WebSphere applications to create event objects, store them and retrieve them later. There are also facilities that pass event objects to registered applications either directly, in the context of the producing (source) application, or indirectly via JMS.
- The ability to fill an event object with additional data needed to correlate events together. This data has to describe the complete execution and business environment to allow different types of users retrieving events to specify queries to satisfy their own data needs.
- A simple event browser (Event browser) for browsing the events that have been stored.

The Event Programming Model exploits the Common Event Infrastructure, a shared component that can be exploited by different types of middleware. The Common Event Infrastructure provides the means to store, update and distribute events to event consumers listening on JMS topics.

**Note:** This information set provides enough information for you to make full use of the Event Programming Model but, if you wish to know more about the Common Event Infrastructure, you can read the developers' guide here: "Common Event Infrastructure Developer's Guide" on page 1173

---

### The Common Event Infrastructure in WebSphere

This topic describes how WebSphere takes advantage of the Common Event Infrastructure to work with application events.

The Common Event Infrastructure provides the runtime environment to persistently store and retrieve events from many different programming environments. See Introduction to the common event infrastructure for more information about the common event infrastructure.

An application creates an event object whenever something happens that either should be recorded for later analysis or which may require additional work to be triggered. This event object describes event type, when the application generated the event and a small number of properties that are relevant to the event.

The application creating the event object is called the event source. The event source passes the event object to the common event infrastructure where data from the event object is aggregated with information from the execution environment and an event catalog before being passed to the Common Event Infrastructure.

The Common Event Infrastructure is a shared component that can either be inside or outside WebSphere Application Server. It provides the means to store, update and distribute events to event consumers listening on JMS topics. Event consumers process events outside the environment of the event source. In fact they are typically processing events from a number of event sources. This has two consequences:

- All event sources need to use a common structure for the event. The accepted standard for such a structure is called the Common Base Event. This defines common fields, the values they can take and the exact meanings of these values.

- The Common Base Event needs to contain all of the information needed by the consumers to understand the event. This includes information about the runtime environment, the business environment and the instance of the application object that created the event. See The Common Base Event model for more information.

## Events and Common Base Events

Understanding the relationship between application events and Common Base Events

An **event** occurs when something significant happens in the IT system. An application processing a new customer order or a failure occurring in a critical part of the system are examples of events.

Information about the event is captured in an event object. This event object describes that the event has occurred, when this happened and any data that is relevant to the event.

The Common Base Event is an XML schema defined as part of the Autonomic Computing initiative IBM Autonomic Computing Initiative. It is used to describe a situation that has occurred in the IT system which may be of interest elsewhere. This includes a precise description of what happened and the environment in which it occurred.

There are two major benefits of adopting the Common Base Event model consistently across systems:

- it saves the expense of creating, documenting and maintaining many record structures
- it makes it possible to manage and process data from a variety of subsystems using common utilities.

In addition to structure, the Common Base Event specification provides guidelines on how to set the attributes and elements of the Common Base Event. This consistency is critical to enabling autonomic correlation of Common Base Events relating to a single situation detected by a number of subsystems.

## The structure of the Common Base Event

The different elements of the Common Base Event, and their structure and values.

A Common Base Event has five logical sections:

1. Attributes for the CommonBaseEvent element. See Attributes for the CommonBaseEvent element.
2. Event classification, describing the situation element. The Common Base Event instance can be one of twelve situation categories. See Event classification elements.
3. Source component identity. For Common Base Events generated within WebSphere, this contains the attributes associated with the sourceComponentId element that describe the IT stack that was executing when the situation occurred. See Common Base Event data describing WebSphere as the source.
4. Context Data Fields. For WebSphere Common Base Events, these fields contain context data used to correlate Common Base Event instances together. See Additional Common Base Event context and application data added by BCDS.
5. Application data. Any additional data supplied by a WebSphere application, and its format. See Additional Common Base Event context and application data added by BCDS.

### Attributes for the CommonBaseEvent element

The meaning of each field in the CommonBaseEvent element of the Common Base Event.

The CommonBaseEvent tag groups the data for a Common Base Event instance together. It has a number of attributes that are common to all instances of Common Base Events. The example below shows the typical use of these attributes in WebSphere.

```
<CommonBaseEvent extensionName="HighValueOrder"
 globalInstanceId="A44486D00A4511D7A5F99C9202ABA117"
 creationTime="2003-05-09T14:28:16.410000-05:00"
 sequenceNumber="0"
```

```

 severity="10"
 priority="50">
 :
 :
</CommonBaseEvent>

```

Each field is described below:

- **extensionName** - this is the type, or class, or the Common Base Event instance. It indicates which extended data elements will be included.
- **globalInstanceid** - a globally unique identifier for the Common Base Event instance.
- **creationTime** - when the Common Base Event instance was created. Its format follows the XML schema dateTime format, which includes a time zone offset and means nanosecond accuracy is possible. However, WebSphere will extract the data and time from the JVM and so is limited to millisecond accuracy. Common Base Events generated within a millisecond can be sequenced automatically.
- **sequenceNumber** - this field counts the number of Common Base Events generated so far within a millisecond interval. It is used by some technologies that need to sequence events within a millisecond time period. The use of the sequenceNumber is optional.
- **severity** - a number between 0 and 70 which describes the impact that the event has had on the creator of the record. Its data type is short. The higher the number the more negative the impact. The predefined severity levels, in order of increasing severity, are:
  - UNKNOWN (0)
  - INFORMATION (10) - used when the event record does not relate to an error.
  - HARMLESS (20) - strictly speaking this is an error but will not cause harm.
  - WARNING (30) - should be used when it's appropriate to let the user decide if action is needed.
  - MINOR (40) - should be used to indicate action is needed, but the situation is not serious at this time.
  - CRITICAL (50) - should be used to indicate action is needed NOW and the scope is broad (perhaps an imminent outage to a critical resource will result).
  - FATAL (60) - should be used to indicate an error occurred, but it's too late to take remedial action. The severity attribute is optional.
- **priority** - this indicates the relative order that the records should be processed in. Its data type is short and valid values are between 0 and 100. Their meanings are:
  - FIRST (100) - This is the highest priority possible. Process this record first.
  - URGENT (90) - This record describes a situation that needs to be resolved or processed as soon as possible.
  - HIGH (70) - This is an important record.
  - MEDIUM (50) - This record is of average importance.
  - LOW (10) - This record does not need to be processed immediately
  - LAST (0) - This is the lowest priority possible. Process this record last

The priority is an optional attribute. Where specified, it takes precedence over the severity. (For example, a record with a priority of 70 and a severity of 40 is processed before a record with a priority of 10 and a severity of 50.)

In addition, components from the WebSphere environment may use these attributes.

- **localInstanceid** - an instance id that is meaningful to the creator and consumer of the Common Base Event but is not necessarily globally unique across all uses of the Common Base Event. This instance id can be up to 128 characters long.
- **msg** - used to document a fully formatted user message.
- **repeatCount** - a counter that is used in conjunction with elapsedTime below.

- **elapsedTime** - the elapsed time documents in seconds the time period over which a situation has repeatedly occurred. The repeatCount attribute documents the number of times the situation occurred within the period. This pair of attributes condenses the number of Common Base Event instances required to describe repeating situations.

The attributes for the CommonBaseEvent tag give some very basic information about a Common Base Event instance. Only the creationTime attribute is required. WebSphere will fill this in automatically if the application does not specify a value.

## Event classification elements

The meaning of each field in the Situation element.

Nested within the CommonBaseEvent element are elements giving more detail about the situation. The first of these is the situation element, a standardized classification of the situation. Classification involves adding fields to the Common Base Event instance that describes why the event was created and how it relates to other activity. The purpose is to simplify event filtering and correlation.

Each Common Base Event instance must contain a valid situation element. This classifies the event into one of twelve situation categories so it can be sorted by common utilities. The categories are as follows:

- StartSituation - a component is starting up
- StopSituation - a component is stopping
- ConnectSituation - a component is connecting to another component
- ConfigureSituation - a component is being configured
- RequestSituation - details the progress of a request issued against a component
- FeatureSituation - indicates the status of a particular feature or service
- DependencySituation - details the availability of a dependent component
- CreateSituation - a component is being created
- DestroySituation - a component is being destroyed
- ReportSituation - details report data from a component
- AvailableSituation - details a component's operational state and availability
- OtherSituation - something else

The Common Base Event specification defines additional sub-elements specific to each category. This is an example of a situation element used by the WebSphere Process Choreographer in an event that documents a new process starting.

```
<situation categoryName="StartSituation">
 <situationType reasoningScope="EXTERNAL" xsi:type="Common Base Event:StartSituation">
 <situationQualifier>START_INITIATED</situationQualifier>
 <successDisposition>SUCCESSFUL</successDisposition>
 </situationType>
</situation>
```

If the application does not supply a situation element, WebSphere adds a default one of category "OtherSituation".

```
<situation categoryName="OtherSituation">
 <situationType reasoningScope="EXTERNAL" xsi:type="Common Base Event:OtherSituation">
 <any>Application Event</any>
 </situationType>
</situation>
```

In addition, WebSphere makes use of extended data elements to define two further classification schemes that reflect the use it makes of the Common Base Event. These are the event domain and the event purpose.



The event domain describes the part of the WebSphere platform stack that the Common Base Event instance comes from. It can take one of the following values:

- **Business** -the Common Base Event instance represents an event in the business world, for example, a late payment by a customer, or a contract being signed.
- **Container** - the Common Base Event instance was created by a WebSphere container as a result of processing events in the application. For example, when a method is invoked, or a Business Process Execution Language (BPEL) activity completes. The control of when a process Common Base Event is generated is defined in deployment descriptors. Knowledge of the application design will allow the reader to extract the business significance of the event.
- **System** - the Common Base Event instance represents an event in the underlying infrastructure, such as when a subsystem is initialized or fails unexpectedly.
- **Debug** - the Common Base Event instance is describing information used for problem determination.

Here is an example of an extended data element showing the event domain:

```
<extendedDataElements name="eventDomain" type="string">
 <values>Business</values>
</extendedDataElements>
```

The **event purpose** describes the purpose of the event within the domain. The event purpose values are as follows:

- **Information** - The Common Base Event instance contains useful information about the operation within the domain.
- **Audit** - The Common Base Event instance records a decision made by someone or something within the domain.
- **Exception** - The Common Base Event instance describes an unexpected event within the domain. This may be an error condition or a situation that needs special processing (for example, the receipt of an extremely valuable customer order).
- **Alert** - The Common Base Event instance records a plea for help by the domain since it may need intervention to resolve a situation.
- **Correlation** - The Common Base Event instance records an association between either a number of events, or a number of data items from the domain.
- **Trace** - The Common Base Event instance is recording a point in the processing of the domain.

This is an example of the event purpose extended data element.

```
<extendedDataElements name="eventPurpose" type="string">
 <values>Audit</values>
</extendedDataElements>
```

**Note:** An **event source category** is an arbitrary name that an application can use to group events. For example, an event source in the category "DomesticClaimsCallCentreAgents" might generate events named "insurance\_claim\_start\_auto" and "insurance\_claim\_start\_home". The Common Base Event may have details of the relevant event source categories active when the Common Base Event instance was created. Notice there may be more than one event source category:

```
<extendedDataElements name="eventSourceCategories" type="string">
 <values>OrderProcessing</values>
 <values>WebSphereProcessChoreographer</values>
</extendedDataElements>
```

## Common Base Event data describing WebSphere as the source

The event data describing WebSphere as the source that is automatically provided in the Common Base Event

## Reference section

By default, Common Base Event instances from WebSphere contain the sourceComponentId and no reporterComponentId. The "component" referred to in this tag is the WebSphere server where the situation occurred. This XML fragment shows the structure of sourceComponentId.

```
<sourceComponentId componentIdType="ProductName"
 component="Platform 5.1 [BASE 5.1.0.1 cf10405.01] [PME 5.1.0 a0405.03]"
 componentType="WebSphereApplicationServer"
 subComponent="Workflow_Engine"
 locationType="Hostname"
 location="9.20.227.85"
 instanceId="alexhome\alexhome\server1"
 executionEnvironment="Windows XP[x86]#5.1"
 processId="1003550944522"
 threadId="Servlet.Engine.Transports : 2" />
```

All of these values are set automatically by the WebSphere runtime unless the caller provides the value explicitly. The attributes are defined as follows:

- **componentIdType** - a fixed value set to "ProductName"
- **component** - the identifiers for the products in the WebSphere platform stack. For example: Platform 5.1 [BASE 5.1.0.1 cf10405.01] [PME 5.1.0 a0405.03]
- **componentType** - a fixed value set to "WebSphereApplicationServer". This value is defined in the Web Application Server Component Types ([http://www.ibm.com/namespaces/autonomic/WebAppServer\\_componentTypes](http://www.ibm.com/namespaces/autonomic/WebAppServer_componentTypes)).
- **subComponent** - the name of the subcomponent within the server that created the event. The values in use are:
  - J2EE\_Application - the default value
  - WebAppServer\_EJB\_Container - generated by a deployment descriptor
  - Workflow\_Engine - generated by the WebSphere Process Choreographer These values are defined in the Web Application Server Component Types ([http://www.ibm.com/namespaces/autonomic/WebAppServer\\_componentTypes](http://www.ibm.com/namespaces/autonomic/WebAppServer_componentTypes)).
- **locationType** - a fixed value set to "Hostname"
- **location** - the hostname of the server (or server region for zOS)
- **instanceId** - the cell name, node name and server name separated by "." characters. It is possible to omit the cell name. It is possible to omit the node name provided the cell name is also omitted. This gives three alternatives:
  - **cell.node.server**
  - **node.server**
  - **server**
- **executionEnvironment** - a string describing the operating system that the application server is running in. This is extracted from the JVM system properties:
  - **os\_name** - Name of the operating system
  - **os\_arch** - Machine architecture
  - **os.version** - Operating system version as follows:

```
String osEnv = new String(System.getProperty("os.name") + "[" +
 System.getProperty("os.arch") + "]"# +
 System.getProperty("os.version"));
```
- **processId** - this is the jvm\_id extracted from the JVM
- **threadId** - this is the identifier for the thread where the Common Base Event instance was created.

## Additional Common Base Event context and application data added by BCDS

The additional context and application data that is added to the Common Base Event by the BCDS.

A Common Base Event instance may contain references to one or more "contexts". An instance of a context gathers - or groups - related types of work together while the application is executing. There are many types of context each managing a different type of work grouping. Some context types, such as transactions and security, are managed by system services that wish to act consistently across a grouping of work. Others are for information purposes only.

The WebSphere Business Context Data Service (BCDS) offers a generic context service to applications. It also adds details of the active BCDS contexts into each Common Base Event instance created so that events from a single instance of a context can be correlated together.

A ContextDataElements element records

- the context **type**
- the context **identifier**

for each active context instance. The example below shows a fragment of a Common Base Event that describes a BCDS context. The WebSphere Process Choreographer has set up an instance of this type of context for each BPEL process.

```
<contextDataElements name="WebSphereProcessChoreographer" type="string">
 <contextValue>OrderCreationTemplate/37471839182</contextValue>
</contextDataElements>
```

Any **application data fields** added to the Common Base Event by an application are saved in extended data elements, as name/value children of the context data. The example below shows three such data fields called customerNo, orderNo, and orderValue.

```
<extendedDataElements name="BusinessTask" type="noValue">
 <children name="customerNo" type="string">
 <values>C03738927</values>
 </children>
 <children name="orderNo" type="string">
 <values>056232-2003-May</values>
 </children>
 <children name="orderValue" type="string">
 <values>1394000</values>
 </children>
</extendedDataElements>
```

## Components you should know about

Event Programming Model components that are important for generating and using events in WebSphere.

There are four components that you should know about to make best use of the Event Programming Model capabilities:

- State observer plugin
- Business Context Data Services (BCDS)
- Common Base Event content handler
- Event browser

### State observer plugin

The state observer plugin processes events triggered from WebSphere Process Choreographer.

The state observer plugin (SOP) accepts events generated from BPEL processes run in WebSphere Process Choreographer and creates events to be passed to the Common Event Infrastructure.

Events will be sent to the state observer plugin if the businessRelevant flag is turned on. This flag can be set for at the process level and also for individual variables and activity instances. See Generating a Common Base Event in BPEL for further information on the businessRelevant flag and working with WebSphere Process Choreographer to generate events.

## Business Context Data Services

BCDS translates between real-world business models and their implementation in IT systems.

The Business Context Data Service (BCDS) makes it possible to define how a running application relates to a conceptual business model. This data provides important content for events generated by the application. The BCDS has two important functions:

- It populates the Common Base Event with relevant data about the business context that caused the event to be generated.
- It manages the data relating to the business model. The data is stored in a temporary area during the lifetime of the process and the BCDS will generate messages if the data is invalid or compromised.

## Common Base Event content handler

Completing event content automatically.

In some situations, you might want some event property data to be automatically set for every event you create. This can be useful if you want to fill in certain standard values that do not change (such as the application name), or if you want to set some properties based on information available from the execution environment (such as creation time or thread information). You might also want to set policies that govern event content according to business rules; for example, you might require that any event with a particular extension name have its severity set to a certain value.

The Common Event Infrastructure provides a content handler mechanism to accomplish this. A content handler is an object that automatically sets the property values of each event based on any arbitrary policies you want to use. The event infrastructure imposes no restrictions on how a content handler modifies event data, so long as the event still conforms to the Common Base Event specification.

The Event Programming Model in WebSphere provides a standard content handler to complete the runtime and context data in the event. You can see the data value mapping in Common Base Event data describing WebSphere as the source.

Common Event Infrastructure also allows you to develop content handlers to implement your own policies. For more information, see [Completing event content automatically](#).

## Event browser

The event browser is a Web application for retrieving events held in the Common Event Infrastructure data store.

Events are saved to a collection of database tables called the event data store. The event browser is a Web application (supplied with the Event Programming Model) that you can use to query and report on the events in the Common Event Infrastructure data store.

The Common Event Infrastructure makes it possible to query the contents of the event data store through a session bean called the Event Access interface: see in the [Common Event Infrastructure information](#), [Creating an event access bean](#).

---

## Starting to use the Event Programming Model

Installing, configuring and designing for the Event Programming Model.

1. Plan how you'll make use of events in your IT systems. See [Planning to use the Event Programming Model](#)
2. Install the Event Programming Model. (The Event Programming Model is installed automatically at the same time as the Common Event Infrastructure.) See [Installing the Event Programming Model](#)
3. Configure services associated with the Event Programming Model. See [Configuring the Event Programming Model using the Administration console](#)

## Planning to use the Event Programming Model

Defining the architecture of your event generation and consumption.

The Common Event Infrastructure provides facilities for generation, propagation, persistence, and consumption of events, but it does not define the events themselves. You should provide the semantics of event management by defining event types and event groups, in the context of an architecture of event sources and event consumers.

1. Identify each **event source** - the application creating the event. The event source passes the event object to the event infrastructure. The event infrastructure will also store the event object in a database for later retrieval. The event infrastructure's role is to pass the event object onto any applications that have expressed an interest in receiving it.
2. Identify each **event consumer** - each application that can make use of the information contained in the event object. Event consumers typically process events from a number of event sources.
3. Identify each **event group**. To recognize that events are a generic concept of use in a variety of situations, the event infrastructure introduces the concept of an event group. This defines the characteristics (property values) that all events of interest to a particular type of consumer would contain. Policies, such as access controls and distribution rules, are assigned to the event groups to customize the behavior of the event infrastructure for each user group.

**Note:** WebSphere supplies a default event group that is defined to include all events. This event group is called "Event groups list" and has a JNDI name of **com/ibm/events/configuration/event-groups/Default**

This figure shows the relationship between these topics:

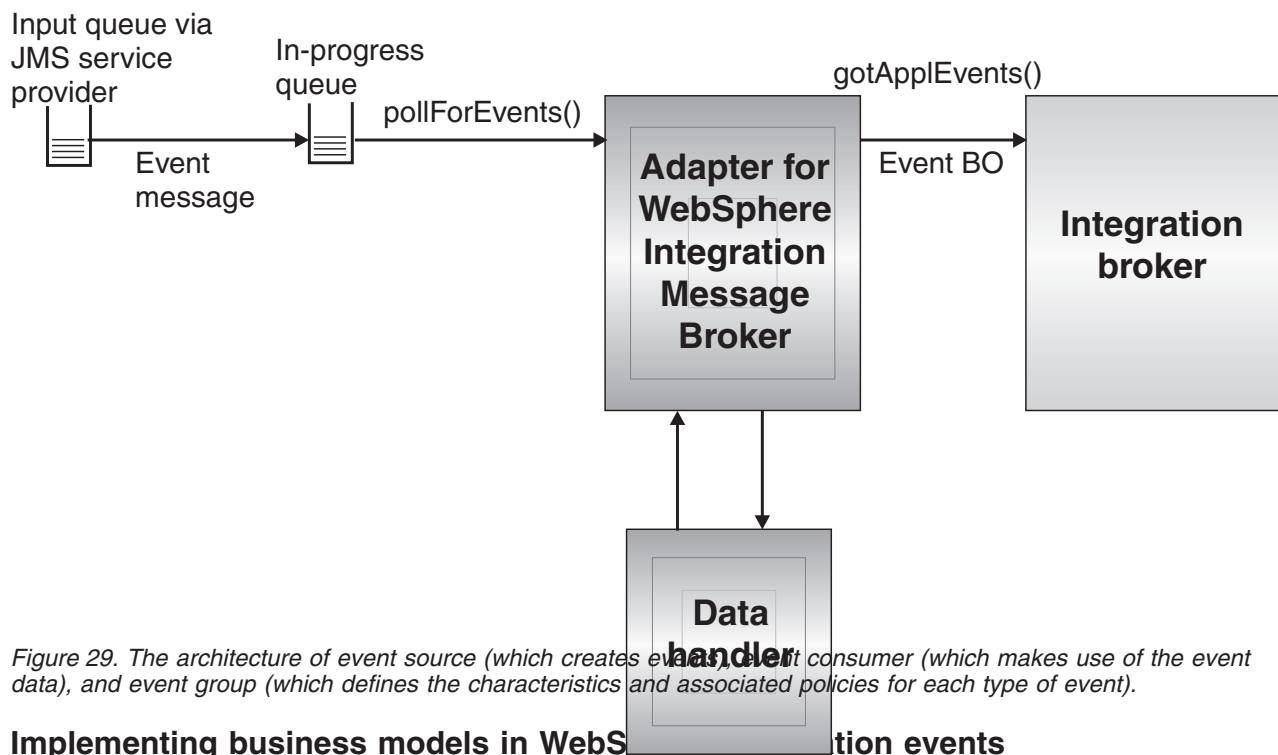


Figure 29. The architecture of event source (which creates event data), event consumer (which makes use of the event data), and event group (which defines the characteristics and associated policies for each type of event).

## Implementing business models in WebSphere Integration events

Things you need to know when designing applications to use the Event Programming Model.

When planning how to use the event infrastructure in your system design, you'll need to understand the business concepts that are relevant, and map them into the appropriate components of your system design.

1. Read about the business model concepts that are relevant to BCDS in Business model concepts.
2. Use the business model concepts to review your WebSphere applications to identify which components implement the concepts you identified. See Using the BCDS in a WebSphere application.

**Business model concepts:**

Things you need to understand before using Business Context Data Services (BCDS).

You need to understand four concepts that appear repeatedly in the many different types of business models.

The first concept is a **business solution**. This describes all of the processing performed by an IT system to support a significant part of the business (for example, order processing). In WebSphere, a solution would include a number of J2EE applications.

There are three more fine-grained concepts:

- The **business task**. This describes the processing to complete a small independent piece of work. (For example, the processing of a single customer order.)
- The **business conversation**. This is the processing performed on behalf of a single person or business partner over a particular period. (For example, the requests made by a single warehouse worker during a single day.)
- The **business data**. This describes the processing that affects a particular collection or instance of some business data. (For example, a single stock item record.)

In a typical IT system, each of these concepts is implemented by one or more components.

Here is an example of a business model that describes part of an order processing solution. Each box is a step in the business task, or process. Part of a warehouse worker’s business conversation involves processing orders (shown in the central lane) and the “Update stock” step updates a single record in the stock database.

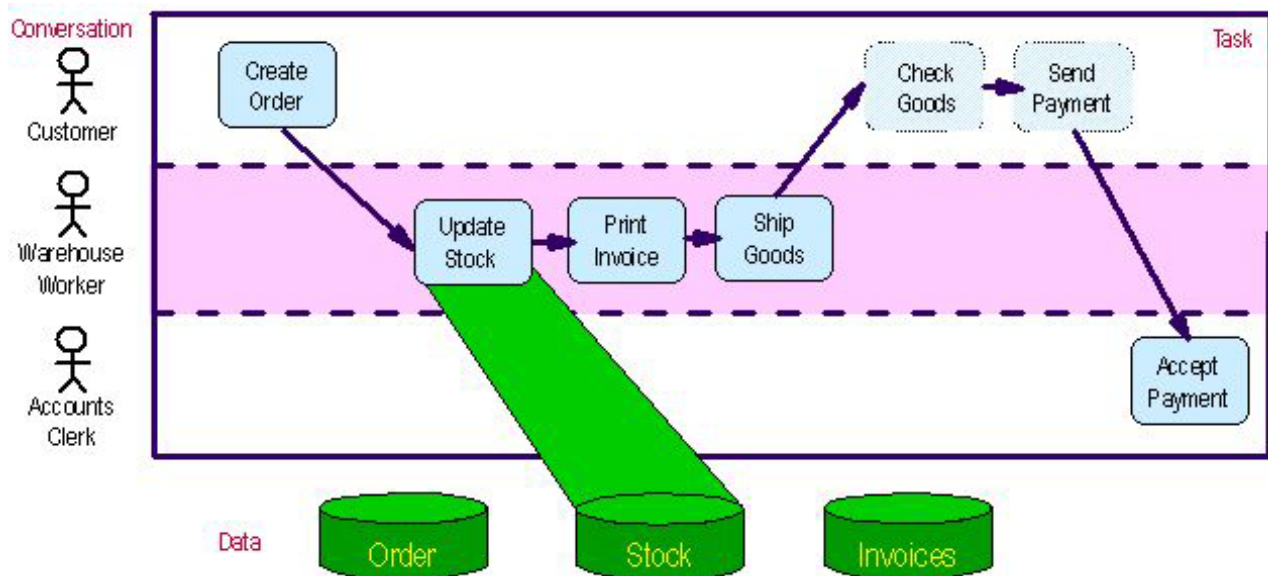


Figure 30. Example business model extract for an order processing solution



During the "Update Stock" activity, you can see that all four of these business concepts (solution, task, conversation and data) intersect. However, during the overall processing of the IT system they change independently of one another.

The BCDS provides the means for an application to create instances of **business context**. Each business context instance has a type (for example, business solution, business conversation, business task or business data), an instance identifier and an optional set of name-value pair properties. The business context instances are associated with the thread so they are available to all components processing a particular request. These components may update and change them at any time. Using the BCDS Java API describes how to use the BCDS to work with the business contexts

**Using the BCDS in a WebSphere application:**

The different ways of using the BCDS in a WebSphere application

1. Work through the business model, or high-level design, and identify the key concepts that are relevant to the application's users. In particular, identify the key user roles, the business tasks that the IT system performs and the key data entities that the business owns. See Business model concepts to read about the business concepts involved: Customer, Warehouse Worker, Accounts Clerk, Fulfil Order Process, Order, Stock, Invoices.
2. Group the business concepts into related types. These groupings will form the basis of your business context types. For example, in Business model concepts the concepts involved are: Customer, Warehouse Worker, Accounts Clerk, Fulfil Order Process, Order, Stock, Invoices (see table below).
3. For each business concept, decide on what the life span of an instance is and how each instance is identified. For example, if there is a business task to "fulfil an order" then there could be a new business task context instance for every order. The order identifier could be the context instance identifier.
4. Review the design of the WebSphere applications to identify which component (or method on a component) implements each of the business concepts identified. You would typically expect servlets, JSPs and portlets to represent the user roles, session beans and/or BPEL process to represent the business tasks and entity beans to represent business entities.
5. Use the BCDS Java programming interface (see Using the BCDS Java API) to update the application components so that a BCDS context instance is active while the application executes the function for an instance of a business concept.

User roles/Conversations	Business Tasks	Business entities/data
Customer	Fulfill order process	Order
Warehouse worker		Stock
Accounts clerk		Invoices

## Installing the Event Programming Model

Installing the Event Programming Model component of WebSphere Business Integration Server Foundation 5.1

You need to have administrator access to the WebSphere Business Integration Server Foundation 5.1 install process.

By choosing to install the Common Event Infrastructure, you also install the Event Programming Model.

To install Common Event Infrastructure and the Event Programming Model in WebSphere, complete the following steps:

1. Use the WebSphere Business Integration Server Foundation 5.1 install program to install the software
  - a. Start the WebSphere Business Integration Server Foundation 5.1 install program



- b. Select the Custom install panel
  - c. Select **Event Infrastructure**, and ensure none of the samples are selected.
  - d. Finish the install program.
2. Configure the Common Event Infrastructure event database.
    - a. Change directory to <WAS\_HOME>/event/dbconfig
    - b. Execute the following command: `config_event_database.bat CloudscapeResponseFile.txt`

The file called CloudscapeResponseFile.txt in this command is a sample response file that will apply to most users. If you have any special circumstances, review the parameters listed in Configuring the event database in the Common Event Infrastructure information, and edit the supplied file.

**Note:** You will need to restart the server after you run the Common Event Infrastructure event database configuration command.

3. Deploy the Common Event Infrastructure EJB.
  - a. Change directory to <WAS\_HOME>/event/application
  - b. Execute the following command: `<WAS_HOME>/bin/wsadmin -f event-application.jacl -action install -earfile event-application.ear -node <nodename> -server <servername>`

**Note:** if your server was started before you issued the above command then you need to start the application from the Admin Console.

You can review the meaning of the parameters listed in the command in Deploying the event infrastructure application in the Common Event Infrastructure information.
4. Install and start the Event Programming Model sample Monitor from the administrative console
  - a. Select **Applications > Install New Application**
  - b. Browse for <WAS root>\eventbrowser\CBEViewer.ear
  - c. Click **Next**.
  - d. Accept the defaults on the next panel titled "Choose default bindings and mappings" and click **Next**.
  - e. Accept the defaults on the first four steps under the panel titled "Allow installation of enterprise applications and modules" by clicking **Next**.
  - f. Review the summary of install options on Step 5 panel, and click **Finish**.
  - g. Check that the message Application CBEViewer is installed successfully appears on the next screen and click **Save to Master Configuration**.
  - h. On the Save panel, click **Save**.

## Configuring the Event Programming Model using the Administration console

You can start and modify the different services you need for the Event Programming Model.

You should have administrator access for the Admin panel

1. Start the Common Event Infrastructure service. See Configuring the Common Event Infrastructure Service.
2. Start the BCDS Service. See Configuring the Business Context Data Service.
3. Start the Application Events Service. See Configuring the Application Events Service.
4. Consider if you need to modify the size of the work area. See Configuring the size of the work area.

## Configuring the Common Event Infrastructure Service

Complete this task to switch the Common Event Infrastructure service on or off for an application server.

The Common Event Infrastructure service provides the ability to process events in WebSphere applications and processes.

You can use the Common Event Infrastructure Service **Startup** property to specify whether or not the service is started automatically for an application server.

To configure the Common Event Infrastructure service **Startup** property for an application server, use the Administrative console to complete the following steps:

1. Start the Administrative console.
2. In the navigation pane, expand **Servers-> Manage Application Servers** This displays a list of the application servers in the content pane.
3. In the Content pane, select the application server that you want to configure. This displays the properties for the application server in the content pane.
4. In the Additional Properties table, select **Common Event Infrastructure service**. This displays the Common Event Infrastructure properties in the content pane.
5. Select or clear the **Startup** property as needed:

**Selected**

[Default] The Common Event Infrastructure service is started when the application server is started. This enables applications that generate events to run on such an application server.

**Cleared**

The Common Event Infrastructure service is not started when the application server is started. Applications that generate events cannot start on such an application server. The BCDS is unaffected. The state observer plugin will not send any events, but will still set business contexts.

Any attempt to start an application that uses Common Event Infrastructure is rejected and a message issued.

If this happens during server startup, the server continues to start without the application.

6. Click **OK**.
7. To save your configuration, click **Save** on the task bar of the Administrative console window.
8. To have the changed configuration take effect, stop then restart the application server.

## Configuring the Business Context Data Service

Complete this task to switch the Business Context Data service on or off for an application server.

The Business Context Data service defines how a running application relates to a conceptual business model. This data provides important content for events generated by applications to achieve end-to-end correlation across a business service, and allows you to define event groups.

You can use the Business Context Data Service **Startup** property to specify whether or not the service is started automatically for an application server.

**Note:** Since there are performance implications in using the BCDS, you may not want to start the BCDS service if:

- event groupings are not important to you, and
- you want to optimize performance potential.

To configure the Business Context Data Service **Startup** property for an application server, use the Administrative console to complete the following steps:

1. Start the Administrative console.
2. In the navigation pane, expand **Servers-> Manage Application Servers** This displays a list of the application servers in the content pane.
3. In the Content pane, select the application server that you want to configure. This displays the properties for the application server in the content pane.
4. In the Additional Properties table, select **Business Context Data Service**. This displays the Business Context Data Service properties in the content pane.

5. Select or clear the **Startup** property as needed:

**Selected**

[Default] The Business Context Data service is started when the application server is started. This enables applications that specify use of Business Context Data in their deployment descriptors to run on such an application server.

**Cleared**

The Business Context Data service is not started when the application server is started. Applications that specify use of Business Context Data in their deployment descriptors cannot start on such an application server.

Any attempt to start an application that uses Business Context Data is rejected and a message issued.

If this happens during server startup, the server continues to start without the application.

6. Click **OK**.
7. To save your configuration, click **Save** on the task bar of the Administrative console window.
8. To have the changed configuration take effect, stop then restart the application server.

### ***Business Context Data Service settings:***

Use this page to administer the runtime properties of the Business Context Data service.

The Business Context Data service defines how a running application relates to a conceptual business model. This data provides important content for events generated by applications to achieve end-to-end correlation across a business service.

To view this administrative console page, click **Servers > Application Servers > *server\_name* > Business Context Data Service**.

#### *Startup:*

Specifies whether the server attempts to start the Business Context Data service automatically.

**Default  
Range**

**Selected  
Selected**

When the application server starts, it attempts to start the Business Context Data service automatically.

**Cleared**

The server does not try to start the Business Context Data service. If applications that run on this server need to use the service, the system administrator must start the service manually or select this property then restart the server.

## **Configuring the Application Events Service**

Complete this task to switch the Application Events service on or off.

The Application Events service provides access to the Common Event Infrastructure for WebSphere applications using the Event Programming Model. It ensures that information about the WebSphere server is automatically included in each event passed to the Event Infrastructure.

You can use the Events Service **Startup** property to specify whether or not the service is started automatically for an application server.

To configure the Events Service **Startup** property for an application server, use the Administrative console to complete the following steps:

1. Start the Administrative console.
2. In the navigation pane, expand **Servers-> Manage Application Servers** This displays a list of the application servers in the content pane.
3. In the Content pane, select the application server that you want to configure. This displays the properties for the application server in the content pane.
4. In the Additional Properties table, select **Application Events Service**. This displays the Events Service properties in the content pane.
5. Select or clear the **Startup** property as needed:
  - Selected**  
 [Default] The Application Events Service is started when the application server is started. This enables applications that specify use of the Common Event Infrastructure in their deployment descriptors to run on such an application server.
  - Cleared**  
 The Application Events Service is not started when the application server is started. Applications that specify use of the Common Event Infrastructure in their deployment descriptors cannot start on such an application server. The state observer plugin will be stopped.  
  
 Any attempt to start an application that uses Events is rejected and a message issued.  
  
 If this happens during server startup, the server continues to start without the application.
6. Click **OK**.
7. Review the JNDI name of the Event Emitter Profile factory that will be used to submit events to the Events Infrastructure. The name provided is of the WebSphere default profile. Unless you have generated an alternative profile, accept the default name.
8. To save your configuration, click **Save** on the task bar of the Administrative console window.
9. To have the changed configuration take effect, stop then restart the application server.

***Application Event service settings:***

Use this page to administer the runtime properties of the Application Event service.

The Application Event service provides access to the Common Event Infrastructure for WebSphere applications and ensures that information about the WebSphere server is automatically included in each event passed to the Common Event Infrastructure.

To view this administrative console page, click **Servers > Application Servers > *server\_name* > Application Events Service**.

*Startup:*

Specifies whether the server attempts to start the Application Event service automatically.

**Default Range**

Selected  
**Selected**

When the application server starts, it attempts to start the Application Event service automatically.

**Cleared**

The server does not try to start the Application Event service. If applications that run on this server need to use the service, the system administrator must start the service manually or select this property then restart the server.

*Event Emitter Factory JNDI Name:*

The JNDI name for the factory that is used to create event emitters.

The Event Submitter first passes the event to its Listener Registry and then (as a Common Base Event) on to the Event Emitter. The event emitter filters the event and if required, passes it on to be stored in the event data store and distributed to the consumers through JMS.

**Data type** String

## Configuring the size of the work area

Configuring the size of the work area to allow greater numbers of contexts.

The BCDS uses the work area to store context data. The default size of the work area is 10,000 and, typically, if there are more than about 15 contexts, this limit is reached with a CORBA.IMP\_LIMIT exception (property group too large). Since the size of a context can vary depending on the data, if you expect to have more than 10 concurrent contexts you should set the size of the work area to expand as necessary.

1. Start the Administrative console.
2. In the navigation pane, expand **Servers-> Manage Application Servers** This displays a list of the application servers in the content pane.
3. In the Content pane, select the application server that you want to configure. This displays the properties for the application server in the content pane.
4. In the Additional Properties table, select **Work Area Service**. This displays the work area properties in the content pane.
5. Select the **Startup** property
6. Set the values of the maxSendSize and the maxReceiveSize fields to **zero**. This has the effect of allowing the work area size to expand with an increasing number of contexts.
7. To save your configuration, click **Save** on the task bar of the Administrative console window.
8. To have the changed configuration take effect, stop then restart the application server.

---

## Generating Common Base Events

Generating Common Base Events from an event source in WebSphere.

An event source is any application that submits events for persistence and distribution to event consumers. The following applications are examples of event sources:

- An adapter or monitor that generates events related to monitored resources
- An application that generates notification events
- An application that forwards events from other sources

An event source must submit valid events conforming to the Common Base Event model. (See The Common Base Event model for more information.) Each event is represented as a Java object.

In general, an event source must follow these high-level steps:

1. Obtain an **emitter**. The emitter is the interface between an event source and the event server. See Emitters and emitter factories.
2. Create an instance of **CommonBaseEvent** (Common Base Event), either by using an event factory through a Java API or from within a process flow described in BPEL, and populate the event with data. See Event creation from WebSphere applications.
3. Submit the event to the event server through the emitter interface. See Send an event.

**Note:** Although it is possible for an application to specify all of the data for a Common Base Event instance, it is only necessary to specify the extensionName. WebSphere is able to fill in the rest of

the details automatically ensuring consistent formatting of this data - but only for events that have occurred within WebSphere. (The Event Programming Model cannot handle events that have been passed to a WebSphere application from another source, like DB2.) For more information about mapping the information in the Common Base Event, see “The structure of the Common Base Event” on page 1144 and, particularly for the information that is filled in by the BCDS, see, “Common Base Event data describing WebSphere as the source” on page 1147 and “Additional Common Base Event context and application data added by BCDS” on page 1148

## Obtain an emitter

How to get an emitter, the first step to creating CBEs.

Before you can obtain an emitter, there must be at least one emitter factory profile configured. See *Configuring the Application Events Service* for information about configuring the emitter factory profile. You can find more information about emitters and emitter factories and the default factory profile in *Emitters and emitter factories*.

To obtain the default emitter, follow these steps:

1. Perform a JNDI lookup specifying the name of the emitter factory you want to use for your emitter. The returned object is the emitter factory
2. Call the `getEmitter()` method of the emitter factory. The returned object is an emitter,

The following example obtains the default emitter:

```
import javax.naming.*
import com.ibm.events.*

Context context = new InitialContext();
EmitterFactory emitterFactory =
 (EmitterFactory) context.lookup("com/ibm/events/configuration/emitter/Default");
Emitter emitter = emitterFactory.getEmitter();
```

## Event creation from WebSphere applications

The different ways to create Common Base Events from within WebSphere applications.

An application creates an event object whenever it performs some processing that could be of interest to an external application. This event object describes event name, when the event happened and some properties that are relevant to the event. There are two different ways of generating events from within WebSphere:

- You can generate an event explicitly from within a WebSphere Java application. See *Using the Java API*.
- You can generate events at different points in application process flows specified in Business Process Flow Language (BPEL). See *Generating a Common Base Event in BPEL*.

Although it is possible for an application to specify all of the data for a Common Base Event instance, it is only necessary to specify the `extensionName`. WebSphere is able to fill in the rest of the details automatically ensuring consistent formatting of this data. For more information about mapping the information in the Common Base Event, see *Common Base Event data describing WebSphere as the source*.

## Using the Java API

Creating a Common Base Event instance from within a Java application.

The example below shows how to create a Common Base Event instance from a J2EE component such as a servlet or EJB.

```
import com.ibm.events.cbe.*;
import javax.naming.Context;

Context ctx = new InitialContext();
EventFactory factory = (EventFactory)ctx.lookup("com/ibm/websphere/events/factory");
```

The code above uses JNDI to obtain the WebSphere-provided event factory; configured with a content handler. (See [Populating events with event context data](#) for more information about content handlers.)

```
CommonBaseEvent event = factory.createCommonBaseEvent("NewOrderReceived");
event.setExtendedDataElement("OrderStatus", "NEW");
event.setExtendedDataElement("CustomerNumber", "C03738927");
```

This code sets three application properties as extended data elements.

The WebSphere event factory will create Common Base Event instances that are capable of automatically filling out the "situation", "sourceComponentId", "contextDataElements" and "extendedDataElements" elements that describe the WebSphere environment whenever the complete() method is called. This may be called explicitly by the application: event.complete();

Alternatively the event emitter (from WebSphere's Event Infrastructure) calls it automatically as part of the sendEvent() method processing.

## Generating a Common Base Event in BPEL

Generating Common Base Events from within BPEL process flows.

Events generated by BPEL are passed to the WebSphere state observer plugin (see [State observer plugin](#) for more information.) The state observer plugin uses the BCDS to convert the events into Common Base Events, the format processed by the Common Event Infrastructure.

WebSphere Process Choreographer can generate six different types of events during the lifetime of a process:

- ProcessTemplate
- ProcessInstance
- ActivityInstance
- ScopeInstance
- Variable
- Link

To activate each type of event, an attribute called wpc:businessRelevant must be set for the corresponding entity in the business model defined by the businessRelevant flag in the BPEL+ extensions supported by WebSphere Process Choreographer. It may be assigned to elements such as activities within the process definition. For further information on the use of the businessRelevant flag to generate events, see the [WebSphere Process Choreographer documentation at "Developing applications for BPEL-based processes"](#).

The example below shows a BPEL process called "purchaseOrderProcess" showing the businessRelevant attribute on the process and activity elements. The effect of this flag is to generate an event each time an instance of the element changes state. It is also possible to call the Common Event Infrastructure from Java snippets and any EJBs called from an activity.

```
<bpel:process name="OrderCreationTemplate" wpc:businessRelevant="yes|no" ...>
...
<bpel:sequence>
<bpel:receive wpc:businessRelevant="yes"
```



```

 bpel:partner="customer"
 bpel:portType="lms:purchaseOrderPT"
 bpel:operation="sendPurchaseOrder"
 bpel:container="PO"/>
...
<bpel:reply wpc:businessRelevant="yes"
 bpel:partner="customer"
 bpel:portType="lms:purchasePT"
 bpel:operation="sendPurchaseOrder"
 bpel:container="Invoice"/>
</bpel:sequence>
</bpel:process>

```

## Populating events with event context data

Completing business context data in an event.

Using the Business Context Data Service, or BCDS, WebSphere's event infrastructure automatically copies the business context data associated with an event into each Common Base Event instance that it creates. Read more about working with the BCDS and how the BCDS and the content handler relate to the Common Base Event in Introduction.

To work with the BCDS, see Using the BCDS Java API.

Using the BCDS Java API includes material on how to develop your own content handler if the WebSphere-supplied version doesn't meet your needs. The Event Mediator is part of the BCDS public API so that application developers can build their own event content handler and corresponding Event Factory to set any additional data in the Common Base Event and still have the business contexts and sourceComponentId element automatically added to the Common Base Event. For more information about the BCDS API, see the Javadoc.

### Introduction

How the Business Context Data Service and the content handler relate to the Common Base Event.

At the heart of the BCDS is the BusinessContextController. This object manages the data for a single business context type and its association with the current thread. This includes the identity of the context and a collection of properties contained in a BusinessContextProperties object. The BusinessContextFactory creates a BusinessContextController. The BusinessContextFactory can also return the types of business contexts active on the thread. The content handler provided by WebSphere adds details of the active business contexts to a Common Base Event

Provided the Event Programming Model event factory is used, the event infrastructure automatically copies the business contexts associated with the current thread into each Common Base Event instance it creates. See Common Base Event data describing WebSphere as the source for details of the data mapping performed when copying into the Common Base Event.

The fragment of a XML below shows the resulting elements in a Common Base Event instance that relates to a single business context instance. Each active context has a <contextDataElements> element with the name attribute set to the context type and the <contextValue> element containing the instance identifier. If there are properties associated with the context, these appear as <children> in an <extendedDataElement> element.

```

<contextDataElements name="BusinessTask" type="string">
 <contextValue>NewOrder-056232030516</contextValue>
</contextDataElements>

<extendedDataElements name="BusinessTask" type="noValue">
 <children name="customerNo" type="string">
 <values>C03738927</values>
 </children>
 <children name="orderNo" type="string">

```

```

 <values>056232-2003-May</values>
 </children>
 <children name="orderValue" type="string">
 <values>1394000</values>
 </children>
</extendedDataElements>

```

In addition, the Event Mediator will set the sourceComponentId element that describes the Event Programming Model. See Common Base Event data describing WebSphere as the source for details of the fields set in the sourceComponentId.

The business context instances associated with the current thread are automatically transported from one WebSphere server to another as calls pass directly between EJBs. However, there is no such automatic mechanism for JMS. If a business context instance needs to pass across JMS in an application message, you should format the active contexts and their properties into an XML document and send as a JMS header property. See Passing business context data across Java Messaging System.

## Using the BCDS Java API

Generating Common Base Events using the BCDS Java API.

The code block below shows how to extract the BusinessContextController for a business task from the current thread and then add some properties to it:

```

import com.ibm.websphere.bcds.*;
import javax.naming.Context;

String factoryName = "com/ibm/websphere/bcds/factory";
String contextType = BusinessContextType.BUSINESS_TASK.getContextType();

Context ctx = new InitialContext();
BusinessContextFactory factory = (BusinessContextFactory)ctx.lookup(factoryName);
BusinessContextController bcc = factory.getController(contextType);
BusinessContextProperties properties = bcc.getProperties();

```

This statement makes the context active in the stack of context instances in the workarea:

```
bcc.setContextId("NewOrder-056232030516");
```

These statements set example properties:

```

properties.setProperty("customerNo", "C03738927");
properties.setProperty("orderNo", "056232-2003-May");
properties.setProperty("orderValue", "1394000");

```

This is how to clear an instance of a business context from the thread:

```
bcc.clearContextId();
```

Since the thread maintains a stack of context instances for each context type, when clearContextId() is called, it restores the context instance for this context type that was on the thread when setContextId() was called.

For more information about the BCDS API, see the Javadoc.

### ***Passing business context data across Java Messaging System:***

How to export business contexts from WebSphere across Java Messaging System (JMS).

The life span of an event can span across several applications and servers. The business context instances associated with the current thread are automatically transported from one WebSphere server to

another as calls pass directly between EJBs. However, there is no such automatic mechanism for JMS. In this case, format the active contexts and their properties into an XML document and send as a JMS header property.

The XML below illustrates the recommended format of an XML document used to import and export business contexts from WebSphere. This XML is typically used to transport the business contexts over JMS in the `com.ibm.websphere.bcds.BusinessContexts` property.

```
<WebSphereBusinessContexts>

 <context typeName="businessSolution" id="Order Processing"/>

 <context typeName="businessConversation" id="JGB-030516">
 <property name="startTime" value="09:00"/>
 <property name="operator" value="Joe G Bloggs"/>
 </context>

 <context typeName="businessTask" id="NewOrder-056232030516">
 <property name="customerNo" value="C03738927"/>
 <property name="orderNo" value="056232-2003-May"/>
 <property name="orderValue" value="1394000"/>
 </context>

 <context typeName="businessData" id="Invoice-134252-030516">
 <property name="creationTime" value="2003-05-09T14:41:45.921000-05:00"/>
 </context>

</WebSphereBusinessContext>
```

The method used to import business contexts from WebSphere is `importContexts`. It sets up the current thread with the business contexts described in the XML document.

The method used to export business contexts from WebSphere is `exportContexts`. It returns an XML document describing the business contexts associated with the current thread. If there are no business contexts associated with the thread, a null is returned.

For more information about the BCDS API, see the Javadoc.

## Send an event

Sending the event to the data store, after it has been completed by BCDS.

To send an event, use the `sendEvent(CommonBaseEvent)` or `sendEvent(CommonBaseEvent, int, int)` method of the `Emitter` interface. You can send an event using the current default emitter settings, or you can specify different options that override the configured defaults for the current event.

To send an event using the current settings, use the `sendEvent(CommonBaseEvent)` method: `String eventId = emitter.sendEvent(event);`

For further information about the actions of the emitter when sending events, see [Sending events](#).

---

## Viewing events with the event browser

This topic describes how to use the event browser to select and review events in the event database.

1. Start the event browser.
  - a. Click on the **Enterprise applications** link in the navigation pane of the WebSphere admin console.
  - b. Click on the tick box beside `CBEViewer`, and click on the **Start** button.
  - c. Check that the status icon is green. Start the `CBEViewer` in your Web browser, using the server name and port 9080, in a URL like `http://<localhost>:9080/cbeviewer/`

2. Specify the events you wish to view. See [Specifying the events to view](#)
3. Select the view of the events returned. See [Working with the events returned](#)

## Specifying the events to view

This task describes how to use the event browser to specify a query.

This task assumes that you have already started the event browser. See [Viewing events with the event browser](#).

1. Specify the calendar period for the report. Enter the start and end dates.
2. Enter the maximum number of records that you want returned.
3. Select the event domains for inclusion, from **Business**, **Process** and **System**. **Note** that if you select none of the domains, then the search will return events for all domains, including events in which no domain is specified.
4. Select from the different types of event purpose, from: **Information**, **Audit**, **Exception**, and **Correlation**. **Note** that if you select none of the purpose types, then the search will return events for all purpose types, including events in which no event purpose is specified.
5. As appropriate, enter your search criteria in the writable fields. You can specify any of: **Solution** name, **Application** name, **Conversation** name, **Task** name, **Data** name, **Process** name, **Server** name, **Component** name, **Event group**, **Priority** and **Severity**.
6. Specify the data store to be searched. The field is a JNDI name, an EJB reference that can be configured in the Admin console. The WebSphere supplied default is `java:/comp/env/events/access`
7. And finally, click on the **Get data** button.

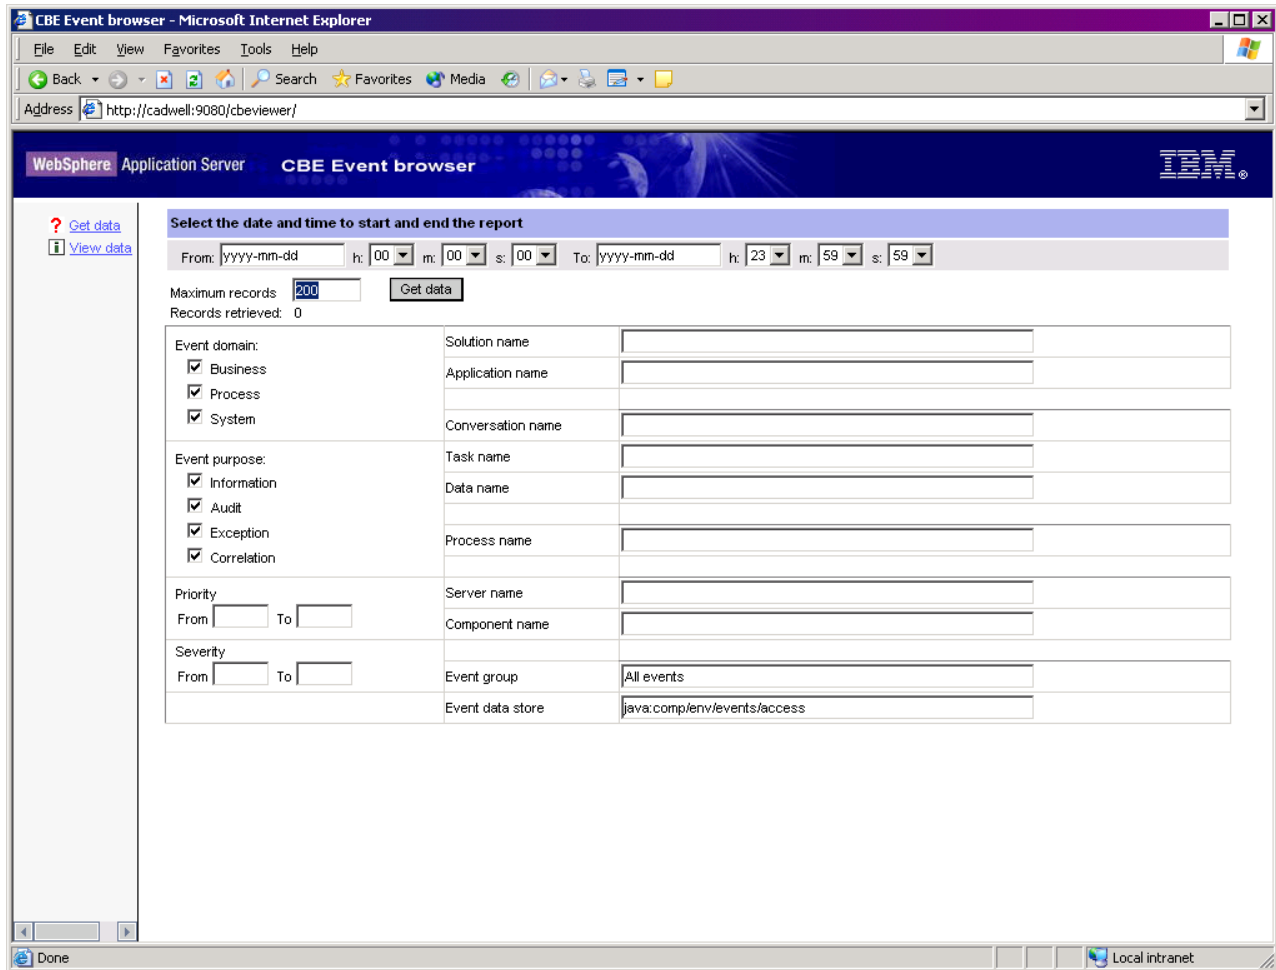


Figure 31. View of browser query screen.

For more information on the origin of these fields, see Origin of event browser fields.

## Origin of event browser fields

The technical details of where each event browser field comes from.

The table below shows the origin of all the fields displayed in the event browser, represented as an XPath expression.

**Note:** The XPath column does not show that all expressions begin with `"/CommonBaseEvent/"`. For example, the full XPath expression to select Creation Time would be `"/CommonBaseEvent/@creationTime"`.

Event Browser field	Position in CommonBaseEvent, as an XPath expression beginning <code>"/CommonBaseEvent/..."</code>
Creation Time	<code>@creationTime</code>
Event Domain	<code>extendedDataElements[@name="eventDomain"]/values</code>
Event Purpose	<code>extendedDataElements[@name="eventPurpose"]/values</code>
Priority	<code>@Priority</code>
Severity	<code>@Severity</code>
Solution Name	<code>contextDataElements[@name="businessSolution"]/contextValue</code>

Application Name	sourceComponentId/@application
Conversation Name	contextDataElements[@name="businessConversation"]/contextValue
Task Name	contextDataElements[@name="businessTask"]/contextValue
Data Name	contextDataElements[@name="businessData"]/contextValue
Process Name	extendedDataElements[@name="websphereProcessChoreographer"] /children[@name="processTemplateName"]/values
Server Name	sourceComponentId/@instanceId
Component Name	extendedDataElements[@name="j2eeComponent"]/children[@name="fullName"]/values
Name	@extensionName
Message	@msg
Instance (on processes tab)	extendedDataElements[@name="websphereProcessChoreographer"] /children[@name="processInstanceId"]/values
name/value properties (Solutions, Conversation, Task and Data tabs)	The complete collection of properties from every event in the collection, related to the selected item. extendedDataElements[@name="NAME"]/children where NAME is related to the currently selected tab (one of "businessSolution", "businessConversation", "businessTask" or "businessData".) And then for each children tag, returns its name and value: children/@name and children/values
name/value at bottom of screen	All values from selected event shown as name/value pairs. For nested properties names are "parentTagName / childTagName"

## Working with the events returned

This task describe show to use the event browser to view the events returned from a query.

This task acts on data returned by a query submitted, as described in Specifying the events to view.

The query returns all the records that meet your criteria. The initial view displayed if you click the **All** tab is shown in the following figure. Click on any event returned to see the event data displayed in the pane at the bottom of the window.

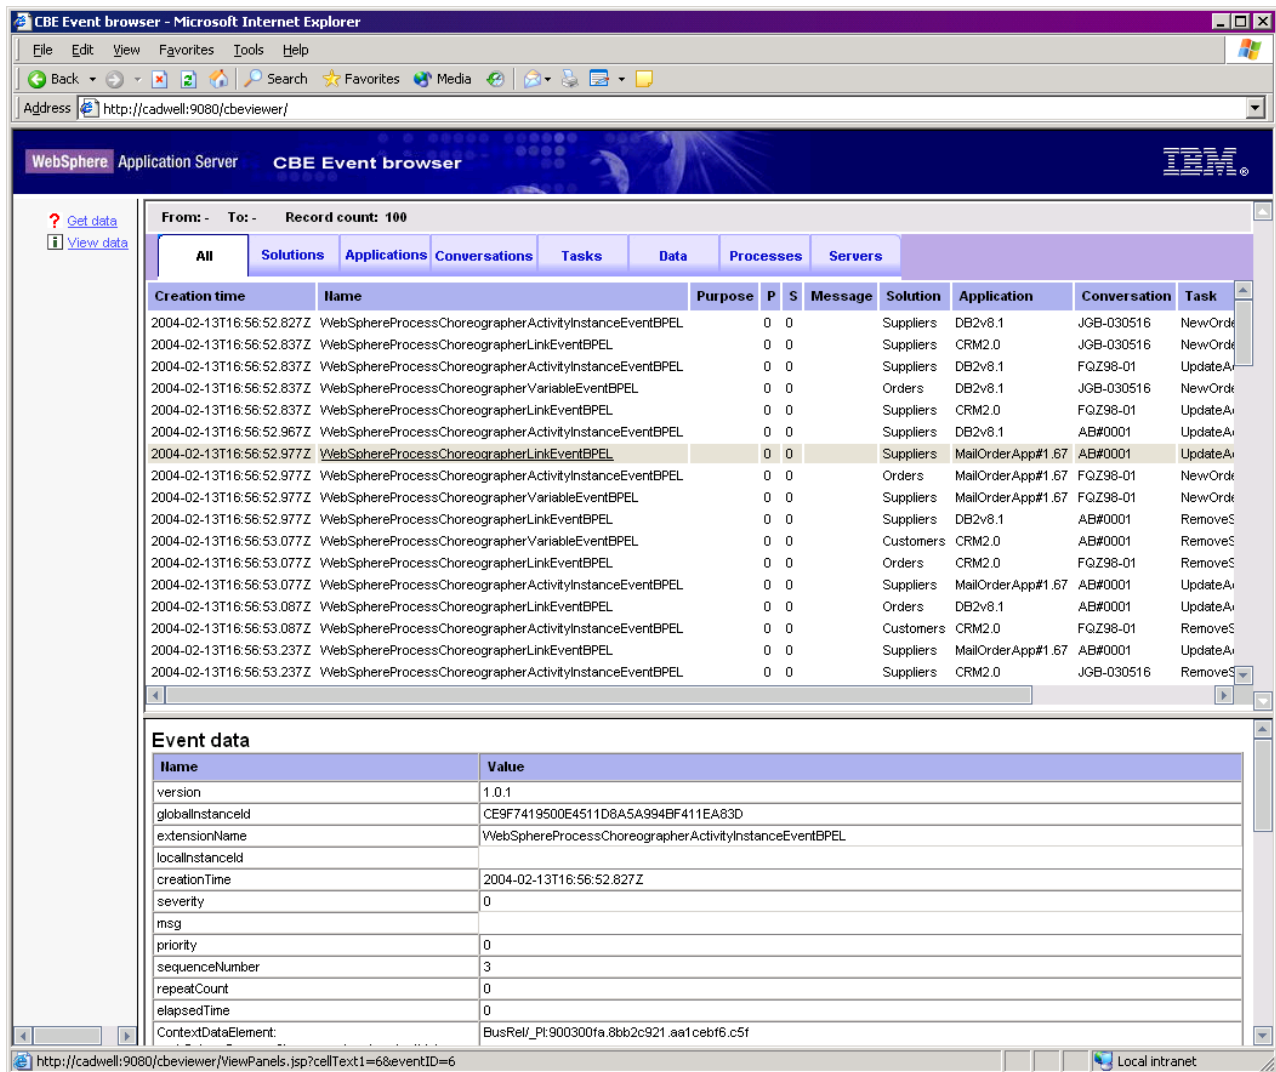


Figure 32. Initial view of event records displayed for an event query.

1. Refine your view of the events returned by clicking on one of the other tabs. These tabs correspond to the selection criteria that you worked with when specifying the query.
2. Further refine your view by selecting from the additional pane that appears when you click on a tab. For instance, if you select the **Solutions** tab, you will be able to select from the whole range of solutions that met your criteria, and selecting one will show you the event data that corresponds to your selection.

**Note:** If you click on the **Processes** tab, you'll be presented with an **Instances** pane when you select a process, to give an extra level of refinement to the search.

If you click on the **Solutions**, **Conversations**, **Tasks** and **Data** tabs, you will see an extra panel that contains any application data fields added to the Common Base Event by an application. They are saved in extended data elements, as name/value children of the context data. See Additional Common Base Event context and application data added by BCDS for more information. For instance, if you click on the **Conversations** tab and select a conversation, the pane will contain name/value pairs of the conversation, extracted from the listed events in that conversation.



---

## Working with WebSphere event infrastructure messages

The meaning of messages from BCDS and state observer plugin

1. If the message code is prefixed by ESOP then it comes from the state observer plugin. You can find the meaning of the messages in State observer plugin messages.
2. If the message code is prefixed by BCDS then it comes from BCDS.. You can find the meaning of the messages in Business context data services messages.

### State observer plugin messages

Messages emitted by the state observer plugin

#### ESOP0001E

The Initial Context for Naming is missing. No lookups are possible. See the nested exception for more information.

##### Explanation

The application is running in an environment where there is no JNDI service. Calls return null.

##### User action

The application needs to be run in a WAS environment where JNDI is available.

#### ESOP0002E

The value of parameter "{0}" is invalid:

##### Explanation

The named parameter has an invalid value. It is rejected. It cannot be null.

##### User action

Correct the calling application and try again.

#### ESOP0003E

The object associated with the BCDS factory JNDI lookup is incorrect. See the nested exception for more information.

##### Explanation

The wrong object has been bound into JNDI.

##### User action

The application needs to be run with the correct object bound into JNDI.

#### ESOP0004E

An exception occurred whilst submitting the event to the Common Event Infrastructure. See the nested exception for more information

**Explanation**

An error occurred with the event infrastructure.

**User action**

Consult the message log.

**ESOP0005E**

An exception occurred whilst obtaining a Common Event Infrastructure emitter. See the nested exception for more information.

**Explanation**

An error occurred with the Common Event Infrastructure.

**User action**

Consult the message log.

**ESOP0006E**

The object associated with the event factory JNDI lookup is incorrect. See the nested exception for more information.

**Explanation**

The wrong object has been bound into JNDI.

**User action**

The application needs to be run with the correct object bound into JNDI.

**ESOP0007E**

The object associated with the emitter factory JNDI lookup is null.

**Explanation**

"null" has been bound into JNDI.

**User action**

Consult the message log

**ESOP0008E**

Message content

**Explanation**

Object associated with the emitter factory JNDI lookup is incorrect. See the nested exception for more information.

**User action**

The application needs to be run with the correct object bound into JNDI.

**ESOP0009E**

An exception occurred whilst obtaining a Common Event Infrastructure emitter factory. See the nested exception for more information.

**Explanation**

An error occurred with the Common Event Infrastructure.

**User action**

Consult the message log

**ESOP0010E**

The object associated with the event factory JNDI lookup is null.

**Explanation**

"null" has been bound into JNDI.

**User action**

Consult the message log.

**ESOP0011E**

An exception occurred whilst obtaining a Common Event Infrastructure event factory. See the nested exception for more information

**Explanation**

An error occurred with the event infrastructure.

**User action**

Consult the message log

**ESOP0012E**

Cannot determine whether the event service is running. The event service is deemed to be running if an event factory can be obtained via JNDI. The object associated with the event factory JNDI lookup is incorrect. See the nested exception for more information

**Explanation**

The event service is deemed to be running if an event factory can be obtained via JNDI. However, the wrong object has been bound into JNDI.

**User action**

The application needs to be run with the correct object bound into JNDI.

## Business context data services messages

Messages issued by the business context data services

### BCDS0001E

No WebSphere work area support is available in this environment.

#### Explanation

The application is running in an environment where the WebSphere work area facility is not available. Calls return null.

#### User action

The application needs to be run in a WAS environment where the work area facility exists.

### BCDS0002E

The Initial Context for Naming is missing. No lookups are possible.

#### Explanation

The application is running in an environment where there is no JNDI service. Calls return null.

#### User action

The application needs to be run in a WAS environment where JNDI is available.

### BCDS0003E

The binding for the WebSphere work area interface was not found. BCDS is not available.

#### Explanation

The application is running in an environment where the WebSphere work area facility is not available. Calls return null.

#### User action

The application needs to be run in a WAS environment where the work area facility has not been disabled.

### BCDS0004E

No active context is associated with the current thread. The action fails.

#### Explanation

The application is unable to access the WebSphere work area facility because it has not been begun on the current thread.

#### User action

Determine how the application came to be executing on a thread where the work area facility has not been begun.

## **BCDS0005E**

The value of property "{0}" cannot be reset or cleared by this process.

### **Explanation**

The WebSphere work area facility only allows a property to be reset or cleared by the process that created it.

### **User action**

Change the application so that properties are reset or cleared only by the originating process.

## **BCDS0006E**

The mode for property "{0}" is read-only. Its value cannot be reset.

### **Explanation**

The WebSphere work area facility does not allow a property whose mode is read-only to be reset.

### **User action**

Check whether the application has corrupted contents of the WebSphere work area that belong to Event Programming Model.

## **BCDS0007E**

The mode for property "{0}" is fixed. Its value cannot be cleared.

### **Explanation**

The WebSphere work area facility does not allow a property whose mode is fixed to be reset.

### **User action**

Check whether the application has inadvertently corrupted contents of the WebSphere work area that belong to WebSphere event infrastructure.

## **BCDS0009E**

The value of parameter "{0}" is invalid.

### **Explanation**

The named parameter has an invalid value. It is rejected.

### **User action**

Correct the application and try again.

## **BCDS0010E**

An error occurred while parsing this imported XML.

**Explanation**

The XML does not conform with the required XSD schema. There may have been a classloader error.

**User action**

Check the XML is valid by reading it into a browser. It may have been corrupted.

**BCDS0011E**

The value for key "{0}" is not serializable. It cannot be returned to the caller.

**Explanation****User action****BCDS0012E**

Identifier "{1}" already exists for business context type "{0}"

**Explanation**

The given identifier is a duplicate for the business context type. It is rejected.

**User action**

Correct the application to use unique business context type identifiers.

---

## Common Event Infrastructure Developer's Guide

Using the Common Event Infrastructure.

- "Introduction to the common event infrastructure"
- "Configuring the common event infrastructure" on page 1175
- "Developing an event source" on page 1180
- "Developing an event consumer" on page 1184
- "Working with events" on page 1188

### Introduction to the common event infrastructure

The common event infrastructure is an embeddable technology intended to provide basic event management services to applications that require those services. This common event infrastructure serves as an integration point for consolidation and persistence of raw events from multiple, heterogeneous sources, and distribution of those events to event consumers. Events are represented using the Common Base Event model, a standard, XML-based format defining the structure of an event. For more information, see The Common Base Event model.

By using this common event infrastructure, diverse products that are not tightly coupled with one another can integrate their management of events, providing an end-to-end view of enterprise resources and correlating events across domain boundaries. For example, events generated by a network monitoring application can be correlated with events generated by a security application. Such correlation can be difficult to achieve when each product uses its own approach to event management.

The common event infrastructure provides facilities for generation, propagation, persistence, and consumption of events, but it does not define the events themselves. Instead, application developers and administrators define event types, event groups, and correlation.

## Components

The common event infrastructure consists of the following major components:

### Common Base Event

The Common Base Event component supports the creation of events and access to their property data. Event sources use the Common Base Event APIs to create new events conforming to the Common Base Event model (see see The Common Base Event model); event consumers use them to read property data from received events. For more information about the Common Base Event component, see Working with .

### Emitter

The emitter component supports the sending of events. After an event source creates an event and populates it with data, it submits the event to an emitter. The emitter validates the event to ensure that it conforms to the Common Base Event specification. If the event is valid, the emitter sends the event to the event server. For more information about emitters, see Developing an event source.

### Event server

The event server is the conduit between event sources and event consumers. The event server receives events submitted to emitters by event sources, and it stores them in a persistent data store. Event consumers can then query the event server to retrieve events. For more information about receiving events from the event server, see Developing an event consumer.

In addition, an application or solution using the event infrastructure might also include the following components (which are not part of the infrastructure itself):

### Event source

An event source is any application that submits events to an emitter.

### Event consumer

An event consumer is any application that receives events from the event server.

## The Common Base Event model

The Common Base Event model is a standard defining a common representation of events that is intended for use by enterprise management and business applications. This standard, developed by the IBM Autonomic Computing Architecture Board, supports encoding of logging, tracing, management, and business events using a common XML-based format, making it possible to correlate different types of events that originate from different applications. For more information about the Common Base Event model, see the Common Base Event specification (*Canonical Situation Data Format: The Common Base Event V1.0.1*). The common event infrastructure currently supports version 1.0.1 of the specification.

The basic concept behind the Common Base Event model is the *situation*. A situation can be anything that happens anywhere in the computing infrastructure, such as a server shutdown, a disk-drive failure, or a failed user login. The Common Base Event model defines a set of standard situation types that accommodate most of the situations that might arise (for example, StartSituation and CreateSituation).

An *event* is a structured notification that reports information related to a situation. An event reports three kinds of information:

- The situation itself (what has happened)
- The identity of the affected component (for example, the server that has shut down)
- The identity of the component that is reporting the situation (which might be the same as the affected component)

The Common Base Event specification defines an event as an XML element containing properties that provide all three kinds of information. These properties are encoded as attributes and subelements of the root element, CommonBaseEvent.



The Common Base Event format is extensible. In addition to the standard event properties, an event can also contain extended data elements, which are application-specific elements that can contain any kind of information relevant to the situation. The *extensionName* attribute can be used to label an event with a classification name, which indicates to applications what sort of extended data elements to expect.

For complete details on the Common Base Event format, see the XML schema included in the Common Base Event specification document, at <ftp://www6.software.ibm.com/software/developer/library/ac-toolkitdg.pdf> .

## Configuring the common event infrastructure

Some configuration is required after the common event infrastructure is installed. This includes deploying the event infrastructure application in WebSphere Application Server and configuring the event database.

In addition, you can perform configuration that affects the runtime behavior of the event infrastructure.

### Post-installation configuration

This set of topics describe the configuration that is required after the common event infrastructure is installed. This includes deploying the event infrastructure application in WebSphere Application Server and configuring the event database.

1. Configuring the event database
2. Deploying the event infrastructure application

#### **Configuring the event database:**

After the event infrastructure application is installed, you must configure the event database and Java Database Connectivity (JDBC) provider. To configure the database, you must first create a response file specifying database parameters, and then use the **config\_event\_database** command to generate and run the required SQL scripts.

*Creating a Cloudscape database response file:* The first step in configuring the event database is creating a response file specifying the database configuration parameters.

**Note:** The database configuration parameters differ depending on the type of database being used. However, currently only the Cloudscape database is supported.

The following parameters are specified in the response file:

#### **SHARE\_DB=[server|node|cell]**

The scope at which the configured database will be shared. This parameter is optional; the default value is `server`.

#### **WAS\_SERVER=*was\_server***

The name of the WebSphere application server where the common event infrastructure application will be installed. This parameter is required only if `SHARE_DB` is set to `server`.

#### **DB\_NAME=*name***

The name of the event database. This parameter is optional; the default is `event`.

#### **DB\_SYSTEM\_DIR=*path***

The path to the directory where the event database will be created. This parameter is optional; the default value is `%WAS_HOME%\event\CloudScapeEventDB`.

#### **JDBC\_PROVIDER=*provider***

The name of the JDBC provider you are using. The value must be the name of a JDBC driver supported by WebSphere Application Server Version 5.1. The following drivers are recommended:

- Cloudscape JDBC Provider
- Cloudscape JDBC Provider (XA)
- Cloudscape Network Server Using Universal JDBC Driver

This parameter is required.

**DB\_TYPE=database**

The type of database being configured. Currently, the only supported value is CLOUDSCAPE. This parameter is required.

**PAGE\_CACHE\_SIZE=pages**

The number of memory pages to use for data caching. This parameter is optional; if it is not specified, the default value for the installed database type is used. (The default value for Cloudscape is 4000.)

**SCHEMA\_NAME=cei**

The schema name used for the event database. This parameter must match the schema name specified in the data store profile, which by default is *cei*. For more information about the data store profile, see Data store profile.

**LOG\_DEVICE=device**

The transaction log device to use. Using a separate device for logs can improve performance, although it complicates backup and recovery. This parameter is optional.

The response file can also include comments, indicated by a number sign (#) character in the first column of any line.

*Sample response file for Cloudscape:* The following is a sample response file for configuring a Cloudscape database:

```
#-----
Sample event database configuration response file for Cloudscape
#-----
SHARE_DB=server
WAS_SERVER = server1
DB_NAME=event
JDBC_PROVIDER="Cloudscape JDBC Provider"
DB_TYPE=CLOUDSCAPE
PAGE_CACHE_SIZE=4000
```

You can also find a sample response file called *CloudscapeResponseFile.txt* in the `%WAS_HOME%\event\dbconfig` directory. You can use this file as a starting point and modify it to create your response file.

*Generating and running the SQL scripts:*

After you have created the response file, you can run the **config\_event\_database** command to generate and run the required SQL scripts. Windows and UNIX versions of this command are provided in the `%WAS_HOME%\event\dbconfig` directory.

On a Windows system, run the following command:

```
config_event_database.bat response_file
```

On a UNIX system, run the following command:

```
config_event_database.sh response_file
```

The *response\_file* parameter is the name of the response file. If the response file is in a different directory, you can specify the complete path to the file.

The **config\_event\_database** command generates the appropriate script in the `%WAS_HOME%\event\dbcscripts\cloudscape` directory. It then automatically runs the script to configure the event database. The following scripts are created:

Database and operating system	Script
-------------------------------	--------

Cloudscape on Windows systems	cr_event_cloudscape.bat
Cloudscape on UNIX systems	cr_event_cloudscape.sh

If you later need to recreate the event database, you can run the appropriate script manually without having to regenerate it. These scripts do not require any parameters.

Similarly, the configuration command generates the appropriate script in the %WAS\_HOME%\event\dsscripts\cloudscape directory and then runs the script to configure the JDBC provider. If you later need to recreate the JDBC provider, you can run the appropriate script:

Database and operating system	Script
Cloudscape on Windows systems	cr_cloudscape_jdbc_provider.bat <i>scope</i> [ <i>server_name</i> ]
Cloudscape on UNIX systems	cr_cloudscape_jdbc_provider.sh <i>scope</i> [ <i>server_name</i> ]

The parameters are as follows:

**scope**

The WebSphere Application Server scope at which you want to create the JDBC provider. The valid values are cell, node, and server.

**server\_name**

The name of the WebSphere server where you want to create the JDBC provider, if **scope** is server. (If **scope** is cell or node, this parameter is not required.)

If an error occurs during database configuration, the created databases are removed, and error messages are written to the log file %WAS\_HOME%\logs\event\event\_db.log.

After you finish configuring the event database, you must restart the WebSphere server.

In addition to configuring the event database, you must also deploy the event infrastructure application in WebSphere Application Server. For more information, see Deploying the event infrastructure application.

**Removing the event database:**

When the event database is created, SQL scripts for database removal are also created and stored in %WAS\_HOME%\event\dbscripts\cloudscape. If you need to remove the event database, you can use the appropriate script:

Database and operating system	Script
Cloudscape on Windows systems	rm_event_cloudscape.bat
Cloudscape on UNIX systems	rm_event_cloudscape.sh

The installation also creates scripts for removing the JDBC provider. To uninstall the JDBC provider, go to the %WAS\_HOME%\event\dsscripts\cloudscape directory and run the appropriate script:

Database and operating system	Script
Cloudscape on Windows systems	rm_cloudscape_jdbc_provider.bat <i>scope</i> [ <i>server_name</i> ]
Cloudscape on UNIX systems	rm_cloudscape_jdbc_provider.sh <i>scope</i> [ <i>server_name</i> ]

The parameters are as follows:

**scope**

The WebSphere Application Server scope from which you want to remove the JDBC provider. The valid values are cell, node, and server.

**server\_name**

The name of the WebSphere server from which you want to remove the JDBC provider, if **scope** is server. (If **scope** is cell or node, this parameter is not required.)

**Note:** The event database can be shared among multiple event servers through the same JDBC provider resource configuration. Therefore, the JDBC provider configuration should be removed only if the associated event database has been uninstalled.

***Deploying the event infrastructure application:***

In addition to configuring the event database, you must also deploy the event infrastructure application in WebSphere Application Server. The installation provides a Jacl script you can use with the WebSphere **wsadmin** command to deploy the application.

To deploy the application on a Microsoft Windows system, go to the %WAS\_HOME%\event\application directory and run the following command (all on one line):

```
wsadmin.bat -f event-application.jacl -action install -earfile event-application.ear
-node node_name -server server_name [-trace]
```

On a UNIX system, go to the \$WAS\_HOME/event/application directory and run the following command (all on one line):

```
wsadmin.sh -f event-application.jacl -action install -earfile event-application.ear
-node node_name -server server_name [-trace]
```

The parameters of this command are as follows:

**node\_name**

The WebSphere Application Server node in which the event infrastructure is installed. You can find out the node name by running the following two commands on the target system:

1. %WAS\_HOME%\bin\setupCmdLine
2. echo %WAS\_NODE%

**server\_name**

The WebSphere server in which the event infrastructure is to be deployed.

If the event-application.ear file is in a different directory, you can specify the complete path to the file. The path should use forward slashes (/) on both Windows and UNIX systems.

The optional **trace** parameter causes additional debugging information to be displayed on the standard output.

***Manually removing the event infrastructure application:***

If you need to manually remove the event infrastructure application from WebSphere Application Server, you can use the **event-application.jacl** script.

To remove the application on a Windows system, go to the %WAS\_HOME%\event\application directory and run the following command (all on one line):

```
wsadmin.bat -f event-application.jacl -action uninstall -node node_name
[-server server_name] [-trace]
```

On a UNIX system, go to the \$WAS\_HOME/event/application directory and run the following command (all on one line):

```
wsadmin.sh -f event-application.jacl -action uninstall -node node_name
[-server server_name] [-trace]
```

The parameters of this command are as follows:

**node\_name**

The WebSphere Application Server node from which you want to remove the event infrastructure application.

**server\_name**

The WebSphere server from which you want to remove the event infrastructure. This parameter is optional; if you do not specify a server, the application is uninstalled from all servers in the node.

The optional **trace** parameter causes additional debugging information to be displayed on the standard output.

## Runtime configuration

The event infrastructure is installed as a set of WebSphere Application Server applications, services, and resources. To customize the event infrastructure, you can configure these resources using the WebSphere Application Server administrative console or the wsadmin tool. (For more information about the administrative console and the wsadmin tool, see "Welcome to System Administration" in the Information Center.)

To see the event infrastructure resources in the administrative console, expand the **Resources** list and click **Common Event Infrastructure Provider**. The types of event infrastructure resources are displayed in the **Additional Properties** list. From here you can navigate to the individual resources of each type (for example, you can view a list of all event group profiles or emitter factory profiles). To change the configuration of a resource, click it in the list and then edit the properties you want to change.

**Note:** Some resources and properties are intended to support functions that are not included in this release. These items are displayed in the administrative console, but they cannot currently be modified.

This chapter describes the event infrastructure resources and the typical configuration tasks you might need to perform. For complete information about these resources and their properties, refer to the administrative console and the HTML file %WAS\_HOME%\web\configDocs\index.html.

**Data store profile:** A data store profile defines properties used by the event data store, which is used to persistently store events received by the event server. A default data store profile is provided; usually, no configuration is necessary for this resource.

**Emitter factory profile:** An emitter factory profile defines properties used for an emitter factory, which event sources use to create emitters. The properties in an emitter factory profile affect the behavior of any emitter that is created using the associated emitter factory. You can use the default emitter factory profile or create additional profiles for your event sources to use. You might want to create an additional emitter factory profile to specify a different transaction mode or synchronous transmission profile. For more information about how these options affect the behavior of the emitter, see Emitters and emitter factories.

To create an emitter factory profile, click **Common Event Infrastructure Provider > Emitter Factory Profile > New**. Specify the properties of the new profile as follows:

**Name** The display name for the profile (this is the name that will be displayed in the list of emitter factory profiles).

**JNDI Name**

The JNDI name for the profile (for example, com/ibm/events/configuration/emitter/Default). This is the JNDI name an event source will use to obtain an emitter factory associated with the profile.

### Preferred Transaction Mode

The transaction mode you want to use. Select the check box if you want each event to be sent in a new transaction; otherwise each event is sent within the current transaction.

### Synchronous Transmission Profile JNDI Name

The JNDI name of the synchronous transmission profile you want to use. This must refer to an existing event bus transmission profile. For example, to use the default synchronous transmission profile, specify `com/ibm/events/configuration/bus-transmission/Default`.

Click **Apply** and then **Save** to save your changes.

**Event bus transmission profile:** An event bus transmission profile defines properties used by emitters to access the event server synchronously using Enterprise JavaBeans (EJB) calls; these profiles are used by emitter factory profiles. A default transmission profile is provided; usually, no configuration is necessary for this resource.

**Event group profile:** An event group profile defines an event group (a logical collection of events). Event groups are used to categorize events according to their content; when querying events from the event server, an event consumer can specify an event group to retrieve only the events in that group.

A default event group profile is provided; this profile defines an event group containing all events. You can create additional event group profiles specifying whatever event criteria are appropriate for your application.

To create an event group profile, click **Common Event Infrastructure Provider > Event Group Profile List > Event groups list > Event Group Profile > New**. Specify the properties of the new profile as follows:

#### Event Group Name

A unique descriptive name for the event group (for example, "Critical" or "Application Server"). This name will be used by event consumers when specifying the event group on a query.

#### Event Selector String

An event selector describing the events you want to include in the event group. An event selector is a regular expression defining a set of events based on their property data and written in a subset of XPath syntax. For more information about writing event selectors, see [Writing event selectors](#).

**Note:** Only one event groups list exists. Do not modify this list or create additional lists.

**Event server profile:** An event server profile defines properties affecting the behavior of the event server. A default event server profile is provided; usually, no configuration is necessary for this resource.

**Filter factory profile:** A filter factory profile defines a set of filter criteria used by event emitters. Currently, filtering is not available, so filter factory profiles cannot be created.

**JMS transmission profile:** A JMS transmission profile defines properties used by emitters to access the event server asynchronously using a Java Messaging Service queue. Currently, asynchronous transmission is not available, so JMS transmission profiles cannot be created.

## Developing an event source

An **event source** is any application that submits events for persistence and distribution to event consumers. The following applications are examples of event sources:

- An adapter or monitor that generates events related to monitored resources
- An application that generates notification events
- An application that forwards events from other sources



An event source is implemented in the Java programming language, using either the Java 2 Platform, Standard Edition (J2SE) or the Java 2 Platform, Enterprise Edition (J2EE). An event source must submit valid events conforming to the Common Base Event model (for more information, see The Common Base Event model). Each event is represented as a Java object.

In general, an event source must follow these high-level steps:

1. Obtain an emitter. The emitter is the interface between an event source and the event server. See *Emitters and emitter factories*.
2. Create an instance of `CommonBaseEvent` and populate the event with data. See *Working with events*.
3. Submit the event to the event server through the emitter interface. See *Sending events*.

## Emitters and emitter factories

An event source does not interact directly with the event server; instead, it interacts with an object called an emitter (an implementation of the `com.ibm.events.emitter.Emitter` interface). An emitter is a local object that provides methods for sending events.

In general, the emitter handles the details of event transmission; the developer of an event source does not need to be concerned about the event server location or the underlying transmission mechanism. Details such as these are governed by the *emitter factory profile*, a configuration object created by an administrator and bound in a Java Naming and Directory Interface (JNDI) namespace as an instance of the `com.ibm.events.emitter.EmitterFactory` interface. An emitter factory profile defines the behavior of an emitter by specifying the following settings (for a complete list, see *Emitter factory profile*):

- The preferred *transaction mode*. This setting specifies whether the emitter attempts to send each event in a new transaction or within the current transaction. An event source can change this setting for a particular emitter, but the profile specifies the default value. (This setting is valid only in a J2EE container; the J2SE platform does not provide transaction controls.)
- The preferred *synchronization mode*. This setting specifies whether events are sent using synchronous or asynchronous transmission. *Synchronous transmission* means that the `sendEvent()` method does not return control to the caller until the event has been processed; *asynchronous transmission* means that the method returns immediately after the event is submitted, and the caller has no further information about event processing. Again, an event source can change this setting for an emitter, but the default value is specified by the profile.

**Note:** Currently, only synchronous event transmission (using EJB calls) is supported.

- The *transmission profiles* to use. A transmission profile is a configuration object that defines a specific transmission mechanism for sending events to the event server. An emitter factory profile can specify two transmission profiles, one for synchronous transmission and one for asynchronous transmission. An event source cannot change the transmission profiles used by an emitter. For more information, see *Event bus transmission profile*.

**Note:** Currently, EJB calls are the only supported transmission mechanism.

An administrator can create multiple emitter factory profiles, each one defining a different emitter configuration. An event source obtains an emitter using the emitter factory associated with an existing emitter factory profile; therefore, all emitters created by a particular emitter factory will have the same default behavior. For more information, see *Obtaining an emitter*.

## Obtaining an emitter

Before you can obtain an emitter, there must be at least one emitter factory profile configured. For each emitter factory profile, an emitter factory is automatically created and is accessible using the JNDI name of the emitter factory profile. For more information, see *Configuring the common event infrastructure*.

To obtain an emitter, follow these steps:



1. Perform a JNDI lookup specifying the name of the emitter factory profile you want to use for your emitter. The returned object is the emitter factory associated with the specified profile. (An event source cannot directly access the profile itself.)
2. Call the *getEmitter()* method of the emitter factory. The returned object is an emitter configured with the profile you specified.

If the emitter factory is unable to obtain an emitter, it throws *EmitterException*.

The following example obtains an emitter configured with the profile *Default*:

```
import javax.naming.*
import com.ibm.events.*

Context context = new InitialContext();
EmitterFactory emitterFactory =
 (EmitterFactory) context.lookup("com/ibm/events/configuration/emitter/Default");
Emitter emitter = emitterFactory.getEmitter();
```

## Sending events

An event source sends events in the form of Java objects. Specifically, each event is an instance of a class implementing the *com.ibm.events.cbe.CommonBaseEvent* interface, which is a Java representation of the Common Base Event specification. For more information, see *The Common Base Event model*.

To create new events, you use an *event factory*, which is an object that returns new instances of *CommonBaseEvent* or of the various specialized classes representing event property data. For more information on event factories and event creation, see *Working with events*.

To send an event, use the *sendEvent(CommonBaseEvent)* or *sendEvent(CommonBaseEvent, int, int)* method of the *Emitter* interface. You can send an event using the current default emitter settings, or you can specify different options that override the configured default settings for the current event.

When you submit an event to an emitter, the following things happen:

1. The emitter calls the *complete()* method of the event, triggering optional content completion. See *Completing event content automatically* for more information.
2. The emitter assigns a sequence number and global instance identifier to any event that does not already have them. It then validates the event to ensure that it conforms to the Common Base Event specification.
3. Finally, if the event is valid, the emitter sends the event to the event server for persistence and distribution to event consumers.

If the event is not valid, or if the emitter encounters a problem when trying to send the event to the event server, an exception is thrown.

### **Using the current emitter settings:**

To send an event using the current settings, use the *sendEvent(CommonBaseEvent)* method:

```
String eventId = emitter.sendEvent(event);
```

In this example, *emitter* is an *Emitter* instance, while *event* is a *CommonBaseEvent* instance. The returned value, *eventId*, is the globally unique identifier of the event (the value of the *globalInstanceId* field of *CommonBaseEvent*). If the event does not have a *globalInstanceId* when you submit it, the emitter assigns one automatically.

### **Overriding the current emitter settings:**

When sending an event, an event source can specify options that override the current transaction mode or synchronization mode (or both) currently configured for the emitter. (The default values for these settings are initially those defined by the emitter factory profile; see Emitter factory profile for more information.)

**Note:** Currently, only synchronous mode is supported.

To override the emitter settings, use the `sendEvent(CommonBaseEvent, int, int)` method:

```
String eventId = emitter.sendEvent(event,
 synchronizationMode,
 transactionMode);
```

The parameters are as follows:

*event*

The event object (an instance of `CommonBaseEvent`) you want to send.

*synchronizationMode*

An integer constant defined by the interface `SynchronizationMode`. (Currently, only synchronous transmission is supported.) This should be one of the following constants:

- `SynchronizationMode.SYNCHRONOUS` (send the event synchronously)
- `SynchronizationMode.DEFAULT` (use the current emitter setting)

*transactionMode*

An integer constant defined by the interface `TransactionMode`:

- `TransactionMode.NEW` (send the event in a new transaction)
- `TransactionMode.SAME` (send the event in the current transaction)
- `TransactionMode.DEFAULT` (use the current emitter setting)

**Note:** An emitter might not support all synchronization and transaction modes. The synchronization modes supported by an emitter are defined by the emitter factory profile, and transactions are supported only in a J2EE container. If you attempt to use a mode that is not supported, the emitter throws `TransactionModeNotSupportedException` or `SynchronizationModeNotSupportedException`, as appropriate.

The following example overrides the emitter setting to send an event in a new transaction, but does not override the synchronization mode:

```
String eventId = sendEvent(event,
 SynchronizationMode.DEFAULT,
 TransactionMode.NEW);
```

Emitter options specified on `sendEvent()` apply only to the single event being sent; no changes are made to the emitter settings, and subsequent event submissions are not affected.

### **Changing the emitter settings:**

An event source can make changes to the transaction mode configured for the emitter. (This setting is initially defined by the emitter factory profile; see Emitter factory profile for more information.)

To change the transaction mode, use the `setTransactionMode(int)` method:

```
emitter.setTransactionMode(transactionMode);
```

The *transactionMode* is an integer constant defined by the interface `TransactionMode`:

- `TransactionMode.NEW` (send the event in a new transaction)
- `TransactionMode.SAME` (send the event in the current transaction)

You can also query the current transaction mode using the `getTransactionMode()` method:

```
int transactionMode = emitter.getTransactionMode();
```

The returned value is an integer corresponding to one of the transaction mode constants (either `TransactionMode.NEW` or `TransactionMode.SAME`).

**Note:** Transactions are supported only in a J2EE container.

## Freeing emitter resources

If your event source has finished sending events with a particular emitter, free the emitter resources using the `close()` method:

```
emitter.close();
```

This method releases all resources being used by the emitter.

## Developing an event consumer

An *event consumer* is any application that receives events from the event server. The event consumer receives events in the form of Java objects; it can then use the `CommonBaseEvent` interface to retrieve event property data.

An event consumer receives events by querying the event access interface. You can query events in either of two ways:

- You can specify a global instance identifier to retrieve a specific single event.
- You can specify an event group and retrieve events associated with that event group. You can optionally refine the query further by specifying an additional event selector, retrieving only events that match both the event group and event selector. For more information, see [Event group profile](#).

In addition to receiving events, an event consumer can also purge old events from the data store.

## Creating an event access bean

The event access interface is implemented as a stateless session bean using the Enterprise JavaBeans architecture. To query the event server using the event access interface, an event source must first create an instance of the event access session bean. The event access bean can be either local or remote.

To create an instance of the event access session bean, use the appropriate home interface (either `EventAccessHome` or `EventAccessLocalHome`):

```
// use home interface to create remote event access bean
InitialContext context = new InitialContext();
Object eventAccessHomeObj = context.lookup("ejb/com/ibm/events/access/EventAccess");
EventAccessHome eventAccessHome = (EventAccessHome)
 PortableRemoteObject.narrow(eventAccessHomeObj,
 EventAccessHome.class);
eventAccess = (EventAccess) eventAccessHome.create();
```

## Querying events by global instance identifier

The Common Base Event specification defines an event property called `globalInstanceId` which can be used as a primary key for event identification. The content of this property is a globally unique identifier, generated either by the application or by the event infrastructure. (Although the Common Base Event specification defines the `globalInstanceId` property as optional, the event emitter automatically assigns an identifier to any event that does not already have one. For more information, see [Sending events](#).)

You can retrieve a specific single event from the event server by querying with the `globalInstanceId` of the event you want to retrieve. This can be useful for testing purposes (to confirm that events are being stored in the event database), or to retrieve an event associated with one that was received previously.

To query an event by global instance identifier, use the `queryEventByGlobalInstanceId()` method of the event access bean:

```
CommonBaseEvent event = eventAccess.queryEventByGlobalInstanceId(eventId);
```

The parameter `eventId` is a string containing the identifier you are searching for. The returned value is an instance of `CommonBaseEvent` representing the event whose `globalInstanceId` matches the string you specified. If no event in the data store matches the specified identifier, the method returns null.

## Querying events by event group

An event can be associated with one or more *event groups*. An event group is a logical grouping of events that match a particular event selector; event groups are defined in the event infrastructure configuration. For more information, see [Event group profile](#).

You can use the event access interface to retrieve events that belong to a specified event group. You can further restrict the query results by specifying an additional event selector. The event access interface provides two methods for querying by event group; the first method returns a limited number of events, while the other returns all events belonging to the specified event group (which might be a large number of events).

To retrieve all of the events belonging to an event group, use the `queryEventsByEventGroup(String, String, boolean, int)` method:

```
CommonBaseEvent[] events = eventAccess.queryEventsByEventGroup(eventGroup,
 eventSelector,
 ascendingOrder,
 maxEvents);
```

The parameters of this method are as follows:

### *eventGroup*

A string containing the name of the event group you want to query events from. This must be the name of an existing event group defined in the event infrastructure configuration.

### *eventSelector*

A string containing an optional event selector that further refines the query. The query will return only events that match both the specified event group and the additional event selector. An event selector is specified in the form of an XPath expression (for more information, see [Writing event selectors](#)). If you do not want to specify an additional event selector, this parameter can be null.

### *ascendingOrder*

A boolean value specifying whether the returned events are to be sorted in ascending or descending order according to the value of the `creationTime` property. If this parameter is *true*, the events are sorted in ascending (chronological) order; if it is *false*, they are sorted in descending (reverse chronological) order.

### *maxEvents*

An integer specifying the maximum number of events you want returned.

To retrieve all of the events in an event group, use the `queryEventsByEventGroup(String, String, boolean)` method, which does not specify *maxEvents*.

**Note:** If the number of matching events exceeds the query threshold defined in the data store profile, `QueryThresholdExceededException` is thrown. The default query threshold is 100 000. For more information, see [Runtime configuration](#).

The following code fragment returns all events that belong to an event group called *critical\_hosts* and whose severity is greater than 30 (warning):

```

CommonBaseEvent[] events =
 eventAccess.queryByEventGroup("critical_hosts",
 "CommonBaseEvent[@severity > 30]",
 true);

```

The following example performs the same query, but specifies that no more than 5000 matching events should be returned:

```

CommonBaseEvent[] events =
 eventAccess.queryByEventGroup("critical_hosts",
 "CommonBaseEvent[@severity>30]",
 true,
 5000);

```

## Purging events from the data store

An event consumer or administrative tool can purge events from the data store using the event access interface. You can purge all events from the data store, or you can limit the purge by specifying event groups, event selectors, or both.

To purge events from the data store, use the `purgeEvents()` method of the event access bean:

```

int purged = eventAccess.purgeEvents(eventGroup,
 eventSelector,
 transactionSize);

```

The parameters are as follows:

### *eventGroup*

A string containing the name of the event group that includes the events you want to purge. This must be the name of an existing event group defined in the event infrastructure configuration. If you do not want to specify an event group, this parameter can be null.

### *eventSelector*

A string containing an optional event selector that identifies the events to purge. An event selector is specified in the form of an XPath expression (for more information, see [Writing event selectors](#)). If you do not want to specify an event selector, this parameter can be null.

### *transactionSize*

A nonzero integer specifying the number of events you want purged in a single database transaction. In most cases, you can use the constant `DEFAULT_PURGE_TRANSACTION_SIZE`, which is defined by the `EventAccess` interface.

The `purgeEvents()` method purges all of the events that match all of the criteria you specify. If *eventGroup* and *eventSelector* are both null, all events in the data store are purged. Events that arrive after the purge operation starts are not purged. The returned value is an integer specifying how many events were purged.

**Note:** If the number of matching events exceeds the maximum purge transaction size defined in the data store profile, `PurgeThresholdExceededException` is thrown and no events are purged. The default maximum purge transaction size is 100 000. For more information, see [Runtime configuration](#).

## Writing event selectors

An event selector is a regular expression that defines a set of events based on their property data (attributes or subelements). For example, an event selector might specify all events from a particular host whose severity is greater than 30 (warning). Event selectors are used to define event groups and query the event server.

Because the Common Base Event specification is based on XML, event selectors are written using a subset of XPath syntax. XPath is a standard language used to identify parts of an XML document; for more information, see the XPath specification at <http://www.w3.org/TR/xpath>.

Currently, an event selector can only specify attributes of the root `CommonBaseEvent` element; subelements and attributes of subelements are not supported. The attributes of `CommonBaseEvent` include the following:

- `version`
- `globalInstanceId`
- `extensionName`
- `creationTime`
- `localInstanceId`
- `severity`
- `priority`
- `msg`
- `sequenceNumber`
- `repeatCount`
- `elapsedTime`

A simple XPath event selector that specifies an attribute value takes the following form:

```
CommonBaseEvent[@attribute = value]
```

For example, the following event selector specifies events with the extension name *ApplicationStarted*:

```
CommonBaseEvent[@extensionName = 'ApplicationStarted']
```

You can also test for the presence of an attribute, regardless of its value:

```
CommonBaseEvent[@severity]
```

You can also perform other relational and equality comparisons using any of the following operators:

- `>` (greater than)
- `<` (less than)
- `=` (equal)
- `!=` (not equal)
- `>=` (greater than or equal)
- `<=` (less than or equal)
- `+` (plus)
- `-` (minus)
- `*` (multiplication)
- `div` (floating point division)
- `mod` (modulo)

**Note:** When used to compare times, these operators perform logical comparisons that recognize time zone differences.

You can form compound logical expressions by using the logical `and` and `or` operators to specify multiple conditions. The following example specifies events created before a particular time and with a severity greater than 30 (warning):

```
CommonBaseEvent[@creationTime < '2003-12-10T12:00:00-05:00' and @severity > 30]
```

When using XPath operators, keep the following rules in mind:

- Always specify logical operators using lowercase letters (for example, `and` rather than `AND`).
- Use spaces to separate operators from the surrounding attribute names and values (`@severity > 30` rather than `@severity>30`).

In addition, you can use the XPath string functions `starts-with`, `contains`, and `string-length`:

```
CommonBaseEvent[contains(@msg, 'disk full')]
```

## Working with events

The common event infrastructure represents events as Java objects. Specifically, each event is an instance of a class implementing the `CommonBaseEvent` interface, which is a Java representation of the Common Base Event specification. (For more information, see [The Common Base Event model](#)). At a high level, the life cycle of an event is as follows:

1. To send an event, an event source creates a new instance of `CommonBaseEvent`, populates it with property data, and then submits it to an emitter. (An optional content completion mechanism, if implemented, can populate events with required property data automatically; for more information, see [Completing event content automatically](#).)
2. The emitter validates the event against the Common Base Event specification. If the event is valid, the emitter sends the event to the event server. For more information about event processing by the emitter, see [Sending events](#).
3. The event server stores the event in a persistent data store.
4. An event consumer submits a query to retrieve the event from the data store, and then reads the event property data. For more information about querying events, see [Developing an event consumer](#).
5. When it is no longer needed, the event is purged from the data store.

The Common Base Event specification, which is based on XML, defines two kinds of event property data:

- Properties represented by simple data types, encoded in XML as attributes of the `CommonBaseEvent` element. These include properties such as *globalInstanceId*, *severity*, and *msg*. The `CommonBaseEvent` Java class represents these values as strings or integers, as appropriate.
- Properties represented by complex data types and encoded in XML as subelements of the `CommonBaseEvent` element. These include properties such as *situation*, *sourceComponentId*, and *extendedDataElements*, each of which has nested properties of its own. These complex types are represented by specialized Java objects; for example, the *sourceComponentId* property is represented by an instance of `ComponentIdentifier`.

The `CommonBaseEvent` interface defines getter and setter methods for each property, as well as helper methods to simplify creation of complex properties. An event source uses the setter methods (or the helper methods) to populate an event with property data before submitting it to an emitter; an event consumer uses the getter methods to retrieve the property data from a received event.

### Creating an event object

To create new events in your event source, you use an *event factory*, which is an object that returns new instances of `CommonBaseEvent` or of the specialized classes representing complex property data types.

There are two ways in which you might access an event factory:

- You can create a new event factory using the *event factory factory*. Use this approach if no appropriate event factory is already available. When you create an event factory, you can optionally specify a content handler to provide automatic content completion. For more information about content completion, see [Completing event content automatically](#).
- You can use an existing event factory that has been bound into a JNDI namespace. Use this approach if an event factory has been provided for you to use; this ensures that any events you create will conform to the appropriate business rules, since the event factory might be configured with a content handler.

#### **Creating a new event factory:**

To create a new event factory, use the event factory factory, implemented as the class `EventFactoryFactory`. This class has no instances; instead, it provides two static methods used to create event factories. The choice of which method to use depends upon whether you want to use a content handler to implement automatic content completion. For more information, see [Completing event content automatically](#).



To create a generic event factory with no content handler, use the `createEventFactory()` static method of `EventFactoryFactory`:

```
EventFactory eventFactory =
 (EventFactory) EventFactoryFactory.createEventFactory();
```

To create an event factory with a content handler, use the `createEventFactory(ContentHandler)` method, specifying the content handler you want to use:

```
EventFactory eventFactory =
 (EventFactory) EventFactoryFactory.createEventFactory(contentHandler);
```

In either case, the returned object is an event factory you can use to create new events.

### **Getting an event factory by JNDI lookup:**

If an event factory has already been provided and bound into a JNDI namespace for your event source to use, perform a standard JNDI lookup to retrieve the event factory:

```
import javax.naming.*
import com.ibm.events.*

Context context = new InitialContext();
EventFactory eventFactory =
 (EventFactory) context.lookup("com/ibm/events/EventFactory");
```

The returned object is the provided event factory; if the event factory is configured with a content handler, an instance of the content handler is also transparently created locally. For more information about content handlers and JNDI, see [Completing event content automatically](#).

### **Populating event data**

After you create an event, you must populate the event with property data. Most event properties are defined as optional by the Common Base Event specification, but four are required:

- *version* (a *string* attribute)
- *creationTime* (an XML *dateTime* attribute)
- *sourceComponentId* (a complex *ComponentIdentification* element)
- *situation* (a complex *Situation* element)

**Note:** The version attribute is defined as optional by the Common Base Event specification, but if it is not specified, the default value 1.0 is assumed. Because the event infrastructure supports only version 1.0.1 of the specification, this value must be specified.

If you try to send an event that is missing any of these properties, the emitter rejects the event and throws a `ValidationException`.

The following code fragment creates an event and populates it with the minimal required property data:

```
CommonBaseEvent event = eventFactory.createCommonBaseEvent();

event.setVersion("1.0.1"); //set version

long currentTime = System.currentTimeMillis(); // get current time
event.setCreationTimeAsLong(currentTime); // and set creationTime

// set sourceComponentId (a complex type)
event.setSourceComponentId("Windows", // application
 "svchost.exe", // component
 "tlntsvr.exe", // subcomponent
 "http://www.ibm.com/namespaces/autonomic/Windows", // componentType
 "win386_svc", // componentIdType
 "9.45.72.138", // location
```

```

 "IPV4" // locationType
);

 // create situation object
 Situation situation = eventFactory.createSituation();

 // set situationType to AvailableSituation (a complex type)
 situation.setAvailableSituation("EXTERNAL", // reasoningScope
 "NOT AVAILABLE", // availabilityDisposition
 "STARTABLE", // operationDisposition
 "FUNCTION_PROCESS"); // processingDisposition

 // set situation
 event.setSituation(situation);

```

This example first uses an event factory to create a new event instance, *event*. First it sets the version property; then it retrieves the current system time and uses the `setCreationTimeAsLong(long)` method to set the value of the *creationTime* property. This could also have been done using the `setCreationTime(String)` method, which sets the creation time using the XML *dateTime* format (for example, "1998-07-29T13:12:00-05:00").

The next required property, *sourceComponentId*, is a complex property represented by an instance of *ComponentIdentification*, which has properties of its own. However, it is not necessary to directly instantiate or interact with this object (although it is possible to do so). Instead, the next statement in the example uses a helper method, `setSourceComponentId()`, to specify the nested properties; the helper method uses these values to create an instance of *ComponentIdentification*, which it then uses to set the value of the *sourceComponentId* property of the event.

Similar helper methods exist for setting other complex properties (for example, `setMsgDataElement()`, `addAssociatedEvents`, and `addExtendedDataElement()`). Many of these methods exist in multiple versions with different signatures, making it possible to specify property values in numerous different ways. Refer to the Javadoc API documentation for complete information on these methods.

The last required property in the example, *Situation*, is another complex property. In this case the situation object must be instantiated directly using an event factory; the example then uses a helper method to set the *situationType* property, which is itself a complex subelement.

In an actual application, a useful event would need to include more information than is shown here, but this is the minimum required by the Common Base Event specification. The event is now valid and can be submitted to an emitter.

## Completing event content automatically

In some situations, you might want some event property data to be automatically set for every event you create. This can be useful if you want to fill in certain standard values that do not change (such as the application name), or if you want to set some properties based on information available from the execution environment (such as creation time or thread information). You might also want to set policies that govern event content according to business rules; for example, you might require that any event with a particular extension name have its severity set to a certain value.

You can accomplish this by creating a *content handler*. A content handler is an object that automatically sets the property values of each event based on any arbitrary policies you want to use. The event infrastructure imposes no restrictions on how a content handler modifies event data, so long as the event still conforms to the Common Base Event specification.

To create a content handler, you must create a new Java class implementing the `com.ibm.events.cbe.ContentHandler` interface. This interface defines a single method called `completeEvent(CommonBaseEvent)`; the parameter is the event whose content is to be completed. In your

implementation of this method, you can use the getter and setter methods of `CommonBaseEvent` to process the event property data in accordance with any policies that apply.

The following example is a simple content handler that automatically sets the extension name of each event:

```
import java.io.Serializable;
import com.ibm.events.cbe.CommonBaseEvent;
import com.ibm.events.cbe.CompletionException;
import com.ibm.events.cbe.ContentHandler;

public class BusinessContentHandler
 implements ContentHandler, Serializable {

 public void completeEvent(CommonBaseEvent event)
 throws CompletionException {
 event.setExtensionName("business");
 }
}
```

To use a content handler, you must associate it with an event factory. To do this, specify the content handler when creating the event factory:

```
EventFactory eventFactory =
 (EventFactory) EventFactoryFactory.createEventFactory(contentHandler);
```

The returned event factory is permanently associated with the specified content handler.

An event source does not need to do anything to enable content completion. If an event factory is associated with a content handler, each event it creates carries with it a reference to that content handler. When the event is submitted to an emitter, the event calls the `completeEvent()` method of the content handler, passing a reference to itself. This ensures that the correct policies are applied to the event after the event source has finished setting event-specific properties, but before the event is validated and processed by the emitter.

**Note:** When an event is transmitted from one process to another, the reference to the content handler is not transmitted with it. This is because content completion relies upon the execution environment where the event originated, and the necessary information might not be available elsewhere. This restriction does not affect calls between applications that are local to one another (for example, a call to an enterprise bean using its local interface).

If you want to ensure that all event sources comply with the same policies, you can create an event factory associated with a content handler (using `EventFactoryFactory`) and then bind the created event factory into a JNDI namespace. Instead of creating their own event factories, event sources can then perform JNDI lookups to access the event factory that already exists, without any knowledge of the content handler. If your business rules later change, you can modify the content handler in one place.

When an event source uses JNDI to retrieve an event factory, the content handler is returned along with the event factory. For this reason, the content handler must be serializable.

## Retrieving event data

When an event source receives an event, it can then use the getter methods of `CommonBaseEvent` to retrieve the event property data. For example, the following code fragment retrieves a single event and then reads the content of the `msg` property.

```
CommonBaseEvent event = eventAccess.queryEventByGlobalInstanceId(eventId);
String eventMessage = event.getMsg();
```

If the property you want to retrieve is a complex property (a subelement of `CommonBaseEvent` in the Common Base Event specification), the returned value is an instance of the specialized class representing

the complex data type. You can then use the getter methods of the returned object to retrieve the property data from that object. For example, the following code fragment retrieves the value of *componentId*, which is a complex property; it then retrieves the content of the nested *component* property, which is a string, to read the name of the source component.

```
CommonBaseEvent event = eventAccess.queryEventByGlobalInstanceId(eventId);
ComponentIdentification componentId = event.getSourceComponentId();
String componentName = componentId.getComponent();
```

## Accessing event metadata

The `com.ibm.events.cbe` package, which provides the classes and interfaces required for working with event objects, is based on the Eclipse Modeling Framework (EMF). EMF is a Java framework used to generate application code based on a structured data model; it also provides interfaces in the generated code that can be used to access metadata describing the data model. (Refer to the Eclipse Modeling Framework documentation at <http://www.eclipse.org/emf> for more information about EMF.)

By using these interfaces, EMF-compatible tools can interact with `CommonBaseEvent` event data without any prior knowledge of the data model or access to the implementation. This makes it possible for development tools to generate code that transfers data from other data models into the `CommonBaseEvent` model. Application developers can then focus on writing code that uses the data rather than code that builds the data.

For example, consider an event source that monitors network events and describes its own data model in terms of EMF. With access to both data models, a development tool could display the fields of the event source data model alongside the fields of the `CommonBaseEvent` data model. A developer could then use a graphical interface to indicate how the fields in the event source model are mapped to fields in the `CommonBaseEvent` model; for example, the field `Workstation.name` in the event source data model might correspond to the field `CommonBaseEvent.sourceComponentId.location` in the `CommonBaseEvent` data model. After this connection is made by the developer, the tool can generate code that handles the transfer of data between the two models.

The following code fragment is a simple example of how a development tool might use EMF interfaces to query information about the `CommonBaseEvent` data model and then use that information to interact with an event instance. This example could be part of a simple event consumer; it iterates through all of the fields of an event instance and, for each one, prints the name and value of the field.

```
// event is a valid CommonBaseEvent instance

// Get list of event instance structural features (fields)
List features = event.eClass().getEAllStructuralFeatures();

// iterate through list; print names and values
for (int i = 0 ; i < features.size() ; i++)
{
 EStructuralFeature feature = (EStructuralFeature)features.get(i);
 Object value = eObj.eGet(feature);
 System.out.println(feature.getName() + ":" + value);
}
```

The `CommonBaseEvent` data model is described in the EMF files `cbe.ecore` and `cbe.genmodel`. These files are included with the common event infrastructure development kit; you can import them into an Eclipse-based development environment and then use EMF to generate code that interacts with `CommonBaseEvent` objects.

---

## Message reference

### CEIC0001W

Event group profiles have been configured with duplicate event group names. Event group names: *event\_group\_names*

#### Explanation

Event group names must be unique in each event group profile. The first duplicated event group profile will be overwritten.

#### Administrator Response

Modify the event group names in the event group profile configuration to ensure that the event group names are unique.

### CEIC0002W

The Common Event Infrastructure service failed to start because the Common Event Infrastructure resources failed to bind into JNDI. Exception message: *binding\_error*

#### Explanation

See message.

#### Administrator Response

Consult the Common Event Infrastructure log files to ensure that the Common Event Infrastructure application is installed successfully. Also, see the exception message for additional information about the cause of the problem.

### CEIC0003W

The Common Event Infrastructure service failed to start because a Common Event Infrastructure object failed to bind into JNDI. Object type: *binding\_object* Exception message: *binding\_object\_error*

#### Explanation

See message.

#### Programmer Response

Consult the Common Event Infrastructure log files to ensure that installation of the Common Event Infrastructure application completed successfully. Also, see the exception message for additional information about the cause of the problem.

### CEIC0004W

The binding in JNDI could not be created. Exception message: *reference\_exception*

#### Explanation

See message.

### **Programmer Response**

Consult the Common Event Infrastructure log files to ensure that the installation of the Common Event Infrastructure application completed successfully. Also, ensure that the configuration settings are correct.

## **CEIC0005W**

The object to bind into JNDI could not be created. Exception message: *reference\_exception*

### **Explanation**

See message.

### **Programmer Response**

Consult the Common Event Infrastructure log files to ensure that the installation of the Common Event Infrastructure application completed successfully. Also, ensure that the configuration settings are correct.

## **CEIC0006I**

The Common Event Infrastructure application is started.

### **Explanation**

No additional information is available for this message.

## **CEIDS0001E**

The specified relational database management system is not supported. Database:*rdbms\_name* Version:  
*rdbms\_version*

### **Explanation**

See message.

### **Administrator Response**

Reconfigure the common event infrastructure data source to use a supported relational database management system.

## **CEIDS0002E**

The relational database management system reported the following error. SQL State: *sql\_state* Vendor code: *vendor\_code* Message: *message*

### **Explanation**

The SQL state is the standard JDBC error code for the reported problem. The vendor code is the database vendor specific error code. The message parameter is the localized error message that is returned by the relational database management system.

### **Administrator Response**

Refer to the appropriate relational database documentation for information about the SQL state, the vendor code, and the error message.

## CEIDS0003E

The mapping of a Common Base Event element or attribute to database tables and columns cannot be found. Attribute: *attribute\_name*

### Explanation

The data store stores metadata information in its database. This metadata information specifies how Common Base Event elements and attributes map to tables and columns in the database. Either an error occurred during the installation of the event database or the metadata table information is corrupted.

### Administrator Response

Reinstall the event database.

## CEIDS0004I

The data store is connecting to the specified relational database management system.  
Database:*rdbms\_name* Version: *rdbms\_version*

### Explanation

No additional information is available for this message.

## CEIDS0005E

The maximum purge transaction size of *purge\_threshold* has been exceeded. The purge operation has been stopped, and no events have been purged.

### Explanation

An event purge operation was attempting to delete a large number of events. Because this action can adversely affect system performance and event throughput, the purge operation was stopped.

### Programmer Response

To reduce the number of events to be purged in a single transaction, do one or both of the following:

- Add additional filters to the purge event selector expression.
- Set the maximum purge transaction size to a lower value.

## CEIDS0006E

The query threshold *query\_threshold* has been exceeded. The query operation has been stopped.

### Explanation

An event query attempted to return a very large number of events. Because this action can adversely affect system performance and event throughput, the query operation was stopped.

### Programmer Response

To reduce the number of events to be returned, do one or more of the following:

- Add additional filters to the event selector expression for the query.
- Specify the maximum number of events to be returned.



- If you have sufficient system resources, increase the query threshold value.

## **CEIDS0007E**

An event with the global instance identifier, *global\_id*, already exists in the database. The event was not stored in the database.

### **Explanation**

Each event instance must have a unique global instance identifier.

### **Programmer Response**

Create the event with a unique global instance identifier.

## **CEIDS0008E**

The WebSphere Application Server connection pooling system returned a database connection that is no longer valid.

### **Explanation**

See message.

### **Programmer Response**

Retry the operation.

## **CEIDS0009E**

A database connection could not be obtained from the connection pool after *maximum\_retries* attempts to obtain a connection.

### **Explanation**

The WebSphere Application Server connection pooling system was not able to return a database connection using the current configuration of the connection pool and the configured number of retries.

### **Administrator Response**

Do one or more of the following:

- Reconfigure the data source, increasing the maximum number of connections. This option improves performance and event throughput.
- Reconfigure the data source, increasing the connection timeout value.
- In the data store profile, increase the number of retries.

## **CEIDS0010I**

The *maxCacheEntries* configuration parameter must be greater than or equal to 0. However, *maximum\_statement\_cache\_size* was specified. The *maxCacheEntries* value has been set to 0.

## Explanation

The `maxCacheEntries` configuration parameter specifies how many translated event selector expressions can be cached by the data store. Caching translated event selector expressions can improve the performance of the event query and the event purge because event selector expressions in the cache do not need to be parsed and translated to SQL.

When the `maxCacheEntries` parameter is set to 0, no translated event selector expressions are cached. Therefore, each event selector expression specified in a call to the event query and the event purge methods must be parsed and translated to valid SQL.

## Administrator Response

In the data store profile, set the `maxCacheEntries` value to 0 or greater.

## CEIDS0011E

The data store cannot find the resource reference in JNDI. Resource reference: *resource\_reference\_name*

## Explanation

The data store uses resource references to obtain the data source that is used to obtain connections to the database and for the object that contains the configuration for the data store. The resource reference information specified at installation is not valid. For example, a JNDI name that is not valid was specified for the resource reference.

If the specified resource reference name is `java:comp/env/DefaultDataStoreProfileReference`, the data store cannot access its configuration data. If the specified resource reference name is `java:comp/env/EventDataSourceReference`, the data store cannot access its data source.

## Administrator Response

In the WebSphere Application Server Administrator's console, reconfigure the specified resource reference to a valid JNDI resource.

## CEIDS0013E

The attribute or element *name* specified in the event selector expression is not supported.

## Explanation

Only attributes in the `CommonBaseEvent` element are supported. The following attributes are supported:

- `globalInstanceid`
- `version`
- `extentionName`
- `localInstanceid`
- `creationTime`
- `severity`
- `msg`
- `priority`
- `sequenceNumber`

## Programmer Response

Specify an event selector expression that contains the supported elements and attributes.

## CEIDS0014E

The XPath location path *locationPath* is not valid, and it could not be translated to a valid SQL expression.

### Explanation

The most likely cause of this problem is specifying a location path that acts on an element or attribute in the `CommonBaseEvent` element that does not map directly to a table column in the event database.

## Programmer Response

Specify a location path that acts on elements and attributes that map directly to a table column in the event database.

The following example shows a valid event selector expression. Although the event selector expression is shown here on multiple lines, you must specify the expression on one line.

```
CommonBaseEvent[messageDataElement/
msgCatalogToken]
```

However, because the `msgCatalogToken` element cannot have data that is mapped directly to a column and table in the event database, it cannot be translated to an equivalent SQL expression.

Instead, use the following event selector expression to select all events from the event database that have at least one message catalog token. Although the event selector expression is shown here on multiple lines, you must specify the expression on one line.

```
CommonBaseEvent[messageDataElement/
msgCatalogToken/@value]
```

## CEIDS0015E

The type of the first argument of the `starts-with` or the `contains` function used in the event selector expression is not supported. Argument: *argument*

### Explanation

The first argument of the `starts-with` or the `contains` function must either be a string constant or a location path that can be translated to a column in the event database. The SQL equivalent to these functions, the `LIKE` function, requires that its first argument be either a table column or a string constant.

## Programmer Response

Specify a string constant or a location path that maps to a table column in the event database for the first argument of the `starts-with` or the `contains` function used in the event selector expression.

The following example shows a valid event selector expression using the `starts-with` function.

```
CommonBaseEvent[starts-with(@msg,'My message')]
```

## CEIDS0016E

The type of the second argument of the starts-with or the contains function used in the event selector expression is not supported. Argument: *argument*

### Explanation

The second argument of the starts-with or the contains function must be a string constant. The SQL equivalent to these functions, the LIKE function, requires that its second argument be a string constant.

### Programmer Response

Specify a string constant for the second argument of the starts-with or the contains function that is used in the event selector expression.

The following example shows a valid event selector expression using the starts-with function.

```
CommonBaseEvent[starts-with(@msg,'My message')]
```

## CEIDS0017E

The first argument of the starts-with or the contains function used in the event selector expression is a constant, but it is not a string constant. Argument: *argument*

### Explanation

If a constant is specified for the first argument of the starts-with or the contains function, the constant must be a string constant. The SQL equivalent to these functions, the LIKE function, requires that its first argument be either a string constant or a table column.

### Programmer Response

Specify a string constant or a location path that maps to a table column in the event database for the first argument of the starts-with or the contains function used in the event selector expression.

The following example shows a valid event selector expression using the starts-with function with a constant as the first argument.

```
CommonBaseEvent[starts-with('My','My message')]
```

## CEIDS0018E

The following division operation is not valid. Operation: *operation*

### Explanation

A zero is the second operand in the specified division operation. Division by zero is not valid.

### Programmer Response

Specify the division operation with a positive numeric value as the second operand.

The following example shows a valid event selector expression using the division operation.

```
CommonBaseEvent[@severity = 10 div 1]
```

## CEIDS0019E

No arguments were specified for the length function used in the event selector expression. Function:  
*function*

### Explanation

You must specify one argument for the length function. The SQL equivalent to the length function requires one argument.

### Programmer Response

Specify one argument for the length function.

The following example shows a valid event selector expression using the length function.

```
CommonBaseEvent[string-length(@msg)=10]
```

## CEIDS0020E

The event selector expression function used in the event selector expression is not supported. Function:  
*function*

### Explanation

The following event selector expression functions are supported:

- string-length()
- contains()
- starts-with()

### Programmer Response

Specify an event selector expression that uses a supported function.

## CEIDS0021E

The XPath expression type, *expression\_type*, is not supported.

### Explanation

The following expression types are supported:

- Constant

```
1,2,'thisString'
```

- Location Path

```
CommonBaseEvent[@severity=10]
```

- Operation

```
@severity+10=20
```

- Function

```
string-length(@msg)
```

## Programmer Response

Specify an XPath expression that uses a supported expression type.

## CEIDS0022E

The event selector expression *expression* is not valid. The event selector parsing tool returned the following message: *message*

### Explanation

The following are possible reasons for this exception:

- A predicate bracket is not closed. For example, `CommonBaseEvent[@severity='10` is missing a closing bracket, and it should be specified as `CommonBaseEvent[@severity='10']`.
- A function has the wrong number of arguments. For example, the `starts-with` and the `contains` function require two arguments.
- A quote on a string constant is not closed. For example, `CommonBaseEvent[@severity='10]` is missing a closing quote, and it should be specified as `CommonBaseEvent[@severity='10']`.
- An open parenthesis is not closed. For example, `CommonBaseEvent[(@severity+10=20]` is missing a closing parentheses, and it should be specified as `CommonBaseEvent[(@severity+10)=20]`.

## Programmer Response

Specify a valid event selector expression.

## CEIDS0023E

The event selector expression *expression* is not valid. The expression includes a step that references the `CommonBaseEvent` element, but it does not contain a predicate.

### Explanation

See message.

## Programmer Response

Specify an event selector expression with at least one predicate associated with the `CommonBaseEvent` element.

For example, the following event selector expression returns SQL that selects all events from the event database.

```
CommonBaseEvent[@globalInstanceId]
```

## CEIDS0024E

The function used in the event selector expression, such as the `starts-with` or the `contains` function, is not valid because it requires two arguments. Function: *function*

### Explanation

The specified function requires two arguments.

### **Programmer Response**

Specify two arguments for the specified function.

The following example shows a valid event selector expression using the starts-with function.

```
CommonBaseEvent[starts-with(@msg,'My message')]
```

## **CEIDS0025E**

The arithmetic operation contains an operand that is not a numeric value. Operation: *operation*

### **Explanation**

Only numeric operands are supported for the following arithmetic operators: +, -, \*, div, mod, <, >, <=, >=

### **Programmer Response**

Specify two numeric values for the operation.

The following example shows a valid event selector expression using the addition operation.

```
CommonBaseEvent[@severity = 10 + 1]
```

## **CEIDS0026E**

The operation contains operands that are not compatible. Operation:*operation*

### **Explanation**

The most likely cause of this problem is using a numeric value for one operand and a string value that cannot be converted to a numeric value for the other operand.

### **Programmer Response**

Specify operands with compatible types.

The following example shows a valid event selector expression using the equality operation.

```
CommonBaseEvent[@severity = '10']
```

Because the string value 10 can be converted to the numeric value 10, the operation is valid and the types are compatible.

## **CEIDS0027E**

An error occurred when writing the CommonBaseEvent any element to the database. The operation was stopped. Error: *io\_write\_error*

### **Explanation**

See message.



### **Administrator Response**

See message.

## **CEIDS0028E**

An error occurred when reading the CommonBaseEvent any element from the database. The operation was stopped. Error: *io\_read\_error*

### **Explanation**

See message.

### **Administrator Response**

See message.

## **CEIDS0029E**

An error occurred when writing a CommonBaseEvent extended data element that has a hexBinary data type to the database. The operation was stopped. Error: *io\_write\_error*

### **Explanation**

See message.

### **Administrator Response**

See message.

## **CEIDS0030E**

An error occurred when reading a CommonBaseEvent extended data element that has a hexBinary data type from the database. The operation was stopped. Error: *io\_read\_error*

### **Explanation**

See message.

### **Administrator Response**

See message.

## **CEIDS0031I**

The configuration value for connection retries, *connection\_retries*, is not valid. The value has been changed to 0.

### **Explanation**

The connection retries configuration setting specifies the number of times the data store attempts to obtain a database connection from the connection pool. If a connection cannot be obtained in the specified number of retries, the data store throws an exception. A value of 0 specifies that the data store does not retry the operation.

### **Administrator Response**

In the data store profile, set the connection retries value to 0 or greater.

## **CEIDS0032I**

The configuration value for query threshold, *query\_threshold*, is not valid. The value has been changed to 100 000.

### **Explanation**

The query threshold configuration setting specifies the maximum number of events that can be returned by the queryEvents methods. This value must be 1 or greater.

### **Administrator Response**

In the data store profile, set the query threshold value to 1 or greater.

## **CEIDS0033I**

The maximum purge transaction size, *purge\_threshold*, is not valid. The value has been changed to 100 000.

### **Explanation**

The maximum purge transaction size configuration setting specifies the maximum number of events that can be deleted in a single purgeEvents call. This value must be 1 or greater.

### **Administrator Response**

In the data store profile, set the maximum purge transaction size to 1 or greater.

## **CEIDS0034E**

The operator used in the the event selector expression operation *operation* is not supported.

### **Explanation**

The following event selector expression operators are supported:

- +, -, \*, div, mod
- =, !=, <, >, <=, >=
- and, or

### **Programmer Response**

Specify an event selector expression operation that uses a supported operator.

## **CEIEI0001E**

The reference obtained from the specified JNDI name is not a reference to an event factory. JNDI name: *JNDI\_name*

**Explanation**

An application attempted to look up an event factory from a JNDI namespace, but the object found in the specified location is not a reference to an event factory.

**Administrator Response**

Contact the owner of the application that is creating the event.

**Programmer Response**

Ensure that the specified JNDI name is a reference to an event factory.

**CEIEI0002E**

The event content handler referenced by this JNDI name cannot be converted back from its binary format. JNDI name: *JNDI\_name* Additional information, if any: *error\_message*

**Explanation**

See message.

**Administrator Response**

If you cannot resolve the problem from the additional information provided, contact IBM Customer Support.

**CEIEI0003E**

The event content handler cannot be converted to binary format. Additional information, if any: *error\_message*

**Explanation**

A problem was encountered when the event content handler was being converted to binary format.

**Administrator Response**

If you cannot resolve the problem from the additional information provided, contact IBM Customer Support.

**CEIEI0005E**

The event identifier in the specified Common Base Event property must start with an alphabetic character. ID: *id* Attribute: *attribute\_name* Property: *class\_name*

**Explanation**

An application has attempted to validate an event that does not conform to version 1.0.1 of the Common Base Event specification.

**Administrator Response**

Contact the owner of the application that is creating the event.

### **Programmer Response**

Ensure that the event conforms to version 1.0.1 of the Common Base Event specification.

## **CEIEI0007E**

The specified extended data element is a type that requires one element inside its array of values, but it contained multiple values. Extended data element: *extended\_data\_element\_name*  
Type: *extended\_data\_element\_type* Number of values: *count*

### **Explanation**

An application has attempted to validate an event that does not conform to version 1.0.1 of the Common Base Event specification.

### **Administrator Response**

Contact the owner of the application that is creating the event.

### **Programmer Response**

Ensure that the event conforms to version 1.0.1 of the Common Base Event specification.

## **CEIEI0008E**

The value in the extended data element cannot be converted to the specified type. Value: *value* Extended data element: *extended\_data\_element\_name* Type: *type*

### **Explanation**

An application has attempted to validate an event that does not conform to version 1.0.1 of the Common Base Event specification. An extended data element inside that event contains a value that does not match its type.

### **Administrator Response**

Contact the owner of the application that is creating the event.

### **Programmer Response**

Ensure that the event conforms to version 1.0.1 of the Common Base Event specification.

## **CEIEI0010E**

The extended data element has a type that is not valid. Extended data element: *extended\_data\_element\_name* Type: *extended\_data\_element\_type*

### **Explanation**

An application has attempted to validate an event that does not conform to version 1.0.1 of the Common Base Event specification.

### **Administrator Response**

Contact the owner of the application that is creating the event.

### **Programmer Response**

Ensure that the event conforms to version 1.0.1 of the Common Base Event specification.

## **CEIEI0011E**

The extended data element contains a null value inside its array of values. Extended data element:  
*extended\_data\_element\_name* Type:*extended\_data\_element\_type*

### **Explanation**

An application has attempted to validate an event that does not conform to version 1.0.1 of the Common Base Event specification.

### **Administrator Response**

Contact the owner of the application that is creating the event.

### **Programmer Response**

Ensure that the event conforms to version 1.0.1 of the Common Base Event specification.

## **CEIEI0015E**

The extended data element contains a string that is not a valid boolean string. Extended data element:  
*extended\_data\_element\_name* Type:*extended\_data\_element\_type* String: *string*

### **Explanation**

An application has attempted to validate an event that does not conform to version 1.0.1 of the Common Base Event specification. Valid values for a boolean string are: true, false, 0, and 1.

### **Administrator Response**

Contact the owner of the application that is creating the event.

### **Programmer Response**

Ensure that the event conforms to version 1.0.1 of the Common Base Event specification.

## **CEIEI0016E**

The extended data element contains a dateTime string that is not valid. Extended data element:  
*extended\_data\_element\_name* Type: *extended\_data\_element\_type* dateTime string: *string*

### **Explanation**

An application has attempted to validate an event that does not conform to version 1.0.1 of the Common Base Event specification. The following example shows a valid dateTime value that specifies 1:20 pm on May the 31st, 1999 for Eastern Standard Time which is 5 hours behind Coordinated Universal Time (UTC).

1999-05-31T13:20:00-05:00

### **Administrator Response**

Contact the owner of the application that is creating the event.

### **Programmer Response**

Ensure that the event conforms to version 1.0.1 of the Common Base Event specification.

## **CEIEI0017E**

The extended data element must contain an even number of characters in the hexValue property.  
Extended data element: *extended\_data\_element\_name* Type: *extended\_data\_element\_type*

### **Explanation**

An application has attempted to validate an event that does not conform to version 1.0.1 of the Common Base Event specification.

### **Administrator Response**

Contact the owner of the application that is creating the event.

### **Programmer Response**

Ensure that the event conforms to version 1.0.1 of the Common Base Event specification.

## **CEIEI0018E**

The extended data element must contain hexadecimal characters in the hexValue property. Extended data element: *extended\_data\_element\_name* Type: *extended\_data\_element\_type* Property value: *value*

### **Explanation**

An application has attempted to validate an event that does not conform to version 1.0.1 of the Common Base Event specification.

### **Administrator Response**

Contact the owner of the application that is creating the event.

### **Programmer Response**

Ensure that the event conforms to version 1.0.1 of the Common Base Event specification.

## **CEIEI0019E**

The CommonBaseEvent/situation/category attribute is not valid. Attribute: *situation\_type*

### **Explanation**

An application has attempted to validate an event that does not conform to version 1.0.1 of the Common Base Event specification.

### **Administrator Response**

Contact the owner of the application that is creating the event.

### **Programmer Response**

Ensure that the event conforms to version 1.0.1 of the Common Base Event specification.

## **CEIEI0020E**

The creationTime attribute for the specified attribute is not valid. creationTime attribute: *creation\_time*  
Attribute: *attribute\_name*

### **Explanation**

An application has attempted to validate an event that does not conform to version 1.0.1 of the Common Base Event specification. The following example shows a valid creationTime value that specifies 1:20 pm on May the 31st, 1999 for Eastern Standard Time which is 5 hours behind Coordinated Universal Time (UTC).

```
1999-05-31T13:20:00-05:00
```

### **Administrator Response**

Contact the owner of the application that is creating the event.

### **Programmer Response**

Ensure that the event conforms to version 1.0.1 of the Common Base Event specification.

## **CEIEI0021E**

The length of an element in the specified attribute exceeds the maximum number of characters. Element: *element*  
Attribute: *attribute\_name* Maximum length: *maximum\_length*

### **Explanation**

An application has attempted to validate an event that does not conform to version 1.0.1 of the Common Base Event specification.

### **Administrator Response**

Contact the owner of the application that is creating the event.

### **Programmer Response**

Ensure that the event conforms to version 1.0.1 of the Common Base Event specification.

## **CEIEI0022E**

The attribute value is outside the valid range. Attribute value: *attribute\_value* Attribute name: *attribute\_name*  
Valid range: *minimum\_value* through *maximum\_value*



**Explanation**

An application has attempted to validate an event that does not conform to version 1.0.1 of the Common Base Event specification.

**Administrator Response**

Contact the owner of the application that is creating the event.

**Programmer Response**

Ensure that the event conforms to version 1.0.1 of the Common Base Event specification.

**CEIEI0023E:**

An attribute is missing from the specified Common Base Event property. Attribute: *attribute\_name*  
Property: *class\_name*

**Explanation**

An application has attempted to validate an event that does not conform to version 1.0.1 of the Common Base Event specification.

**Administrator Response**

Contact the owner of the application that is creating the event.

**Programmer Response**

Ensure that the event conforms to version 1.0.1 of the Common Base Event specification.

**CEIEI0024E**

The attributes are mutually exclusive in the specified Common Base Event property. Attribute:  
*attribute\_name* Attribute: *attribute\_name* Property: *class\_name*

**Explanation**

An application has attempted to validate an event that does not conform to version 1.0.1 of the Common Base Event specification.

**Administrator Response**

Contact the owner of the application that is creating the event.

**Programmer Response**

Ensure that the event conforms to version 1.0.1 of the Common Base Event specification.

**CEIEI0025E**

The attribute value in the specified Common Base Event property is not valid. Attribute value: *value*  
Attribute: *attribute\_name* Property: *class\_name*

### **Explanation**

An application has attempted to validate an event that does not conform to version 1.0.1 of the Common Base Event specification. The value must be 1.0.1.

### **Administrator Response**

Contact the owner of the application that is creating the event.

### **Programmer Response**

Ensure that the event conforms to version 1.0.1 of the Common Base Event specification.

## **CEIEI0026E**

The length of the value in the specified Common Base Event property is outside the valid range of 32 to 64 characters. Value: *value* Attribute: *attribute\_name* Property: *class\_name*

### **Explanation**

An application has attempted to validate an event that does not conform to version 1.0.1 of the Common Base Event specification.

### **Administrator Response**

Contact the owner of the application that is creating the event.

### **Programmer Response**

Ensure that the event conforms to version 1.0.1 of the Common Base Event specification.

## **CEIEI0027E**

The length of an attribute in the specified Common Base Event property exceeds the maximum number of characters. Attribute length: *attribute\_length* Attribute: *attribute\_name* Property: *class\_name* Maximum number of characters: *maximum\_number*

### **Explanation**

An application has attempted to validate an event that does not conform to version 1.0.1 of the Common Base Event specification.

### **Administrator Response**

Contact the owner of the application that is creating the event.

### **Programmer Response**

Ensure that the event conforms to version 1.0.1 of the Common Base Event specification.

## **CEIEI0028E**

The length of an attribute value in the specified Common Base Event property is outside the valid range. Attribute length: *attribute\_length* Attribute: *attribute\_name* Property: *class\_name* Valid number of characters: *minimum\_maximum\_length*

**Explanation**

An application has attempted to validate an event that does not conform to version 1.0.1 of the Common Base Event specification.

**Administrator Response**

Contact the owner of the application that is creating the event.

**Programmer Response**

Ensure that the event conforms to version 1.0.1 of the Common Base Event specification.

**CEIEI0029E**

The attributes in the specified Common Base Event property must either all be set or all be unset.  
Attributes: *attribute\_list* Property: *class\_name*

**Explanation**

An application has attempted to validate an event that does not conform to version 1.0.1 of the Common Base Event specification.

**Administrator Response**

Contact the owner of the application that is creating the event.

**Programmer Response**

Ensure that the event conforms to version 1.0.1 of the Common Base Event specification.

**CEIEI0030E**

When the value of the following attribute in the specified Common Base Event property is equal to *attribute\_value*, the specified list of attributes must be unset. Property: *class\_name* Attribute: *attribute\_name* Attribute list: *attribute\_list*

**Explanation**

An application has attempted to validate an event that does not conform to version 1.0.1 of the Common Base Event specification.

**Administrator Response**

Contact the owner of the application that is creating the event.

**Programmer Response**

Ensure that the event conforms to version 1.0.1 of the Common Base Event specification.

**CEIEI0031E**

The situation type is not valid for the specified situation category. Situation type: *situation\_type* Situation category: *situation\_category*.

**Explanation**

An application has attempted to validate an event that does not conform to version 1.0.1 of the Common Base Event specification.

**Administrator Response**

Contact the owner of the application that is creating the event.

**Programmer Response**

Ensure that the event conforms to version 1.0.1 of the Common Base Event specification.

**CEIEM0001E**

The emitter failed to initialize because the configured synchronization mode *mode* is not valid.

**Explanation**

See message.

**Administrator Response**

Ensure that the synchronization mode specified in the emitter factory profile is valid.

**CEIEM0002E**

The emitter failed to start because the configured synchronization mode is not valid. A sender has not been configured for the synchronization mode.

**Explanation**

See message.

**Administrator Response**

Ensure that the synchronous transmission profile specified in the emitter factory profile is valid.

**CEIEM0003E**

The filter factory *filterFactory* failed to create a filter instance.

**Explanation**

See message.

**Administrator Response**

See the exception message that follows this message for additional information about the cause of the problem.

**CEIEM0004E**

The filter factory *filterFactory* failed to create a filter instance.

**Explanation**

See message.

**Administrator Response**

See the exception message that follows this message for additional information about the cause of the problem.

**CEIEM0005E**

A null event was passed to the `CommonBaseEvent` instance.

**Explanation**

The event parameter to the `Emitter.sendEvent(CommonBaseEvent event)` method must be a valid instance of `com.ibm.events.cbe.CommonBaseEvent`.

**Administrator Response**

Change the call to `Emitter.sendEvent(CommonBaseEvent)` to pass a valid `com.ibm.events.cbe.CommonBaseEvent` instance.

**CEIEM0006E**

The emitter failed to filter the event. Event: *event* Filter: *filter*

**Explanation**

See message.

**Administrator Response**

See the exception message that follows this message for additional information about the cause of the problem.

**CEIEM0007E**

The emitter failed to obtain the filter metadata from the filter factory.

**Explanation**

See message.

**Administrator Response**

See the exception message that follows this message for additional information about the cause of the problem.

**CEIEM0008E**

The filter failed to close. Resources held by the filter instance might not have been released. Filter: *filter*

**Explanation**

See message.

**Programmer Response**

See the exception message that follows this message for additional information about the cause of the problem.

**CEIEM0014E**

The emitter does not support the specified transaction mode *mode*.

**Explanation**

See message.

**Administrator Response**

Ensure that a valid transaction mode is configured.

**CEIEM0015E**

The emitter does not support the specified synchronization mode *mode*.

**Explanation**

See message.

**Administrator Response**

Ensure that the synchronization mode specified in the emitter factory profile is valid.

**CEIEM0020E:**

The emitter failed to initialize because the event bus could not be found. The event bus JNDI name is bound to null. JNDI name: *JNDI\_name* Namespace: *namespace*

**Explanation**

See message.

**Administrator Response**

Ensure that the JNDI name of the event bus as specified in the synchronous transmission profile is correct.

**CEIEM0021E**

The emitter failed to initialize because the event bus could not be found. The event bus JNDI name resolves to an instance of *class*. This name must resolve to an instance of *remoteEventBusClassName* or *localEventBusClassName*.

**Explanation**

See message.

**Administrator Response**

Ensure that the JNDI name of the event bus as specified in the synchronous transmission profile is correct.

**CEIEM0022E**

The emitter failed to initialize because the event bus could not be found. The event bus JNDI name defined in the event bus transmission profile does not exist. JNDI name: *JNDI\_name*

**Explanation**

See message.

**Administrator Response**

Ensure that the event bus JNDI name attribute of the event bus transmission profile is correct.

**CEIEM0023E**

The emitter failed to start because of a failure in the synchronous transmission sender. An exception occurred during initialization of the event bus sender. Exception: *exception*

**Explanation**

See message.

**Administrator Response**

See the exception message that follows this message for additional information about the cause of the problem.

**CEIEM0024E**

The emitter was unable to close due to a failure in the synchronous transmission sender. An exception occurred when attempting to close the emitter. A call to the *localOrRemote* event bus EJB `ejbRemove()` method failed. Exception: *exception*

**Explanation****Administrator Response**

See the exception message that follows this message to determine the exact cause of the failure.

**CEIEM0025E**

The *localOrRemote* event bus EJB on the event server failed during event processing, and the emitter failed to send the event to the event server. Event: *event* Transaction mode: *transaction\_mode*



**Explanation**

Either the server is unavailable or is not configured correctly, or an unexpected runtime error has occurred.

**Administrator Response**

See the exception message that follows this message for additional information about the cause of the problem.

**CEIEM0030E**

The lookup of a JNDI name failed. JNDI name: *JNDI\_name*

**Explanation**

See message.

**Administrator Response**

See the exception message that follows this message for additional information about the cause of the problem.

**CEIEM0031E**

The lookup of a JMS queue connection factory failed because the JNDI name defined in the emitter profile is not bound. JNDI name: *JNDI\_name*

**Explanation**

See message.

**Administrator Response**

Ensure that the configuration of the JMS queue connection factory asynchronous transmission profile for this event emitter is correct.

**CEIEM0032E**

The lookup of a JMS queue connection factory failed because the JNDI name does not resolve to an instance of `javax.jms.QueueConnectionFactory`. JNDI name: *JNDI\_name*

**Explanation**

See message.

**Administrator Response**

Ensure that the configuration of the JMS queue connection factory asynchronous transmission profile for this event emitter is correct.

**CEIEM0033E**

The lookup of a JMS queue connection factory failed because the JNDI name was bound to null instead of a valid `javax.jms.QueueConnectionFactory` instance. JNDI name: *JNDI\_name* failed.

**Explanation**

See message.

**Administrator Response**

Ensure that the configuration of the JMS queue connection factory asynchronous transmission profile for this event emitter is correct.

**CEIEM0034E**

The lookup of a JMS queue failed because the JNDI name defined in the emitter profile is not bound.  
JNDI name: *JNDI\_name*

**Explanation**

See message.

**Administrator Response**

Ensure that the configuration of the JMS queue and asynchronous transmission profile for this event emitter is correct.

**CEIEM0035E**

The lookup of a JMS queue failed because the JNDI name does not resolve to an instance of `javax.jms.Queue`. JNDI name: *JNDI\_name*

**Explanation**

See message.

**Administrator Response**

Ensure that the configuration of the JMS queue and asynchronous transmission profile for this event emitter is correct.

**CEIEM0036E**

The lookup of a JMS queue failed because the JNDI name was bound to null instead of a valid `javax.jms.Queue` instance. JNDI name: *JNDI\_name*

**Explanation**

See message.

**Administrator Response**

Ensure that the configuration of the JMS queue and asynchronous transmission profile for this event emitter is correct.

**CEIES0001E**

The event distribution received a null event when attempting to create the event notification.

**Explanation**

See message.

**Programmer Response**

Specify an event when creating the event notification.

**CEIES0002E**

The event distribution could not validate the event selector for the event group. Event selector: *event\_selector* Event group: *event\_group\_name* Exception message: *exception\_message*

**Explanation**

The event selector used by event distribution does not contain a valid XPath expression for filtering events.

**Administrator Response**

Ensure that the event group configuration for the event selector contains a valid XPath expression for filtering events.

**CEIES0003E**

The event distribution failed to initialize an event group because the specified destination could not be found in JNDI. Event group name: *event\_group\_name* Factory JNDI name: *destination\_factory\_jndi\_name* JNDI name: *destination\_jndi\_name* Exception message: *exception\_message*

**Explanation**

See message.

**Administrator Response**

Ensure that the event group definition contains a valid topic connection factory and topic or queue connection factory and queue as referenced in JNDI.

**CEIES0004E**

The event distribution failed to connect to the JMS destination with the specified JNDI name. Event group name: *event\_group\_name* Factory JNDI name: *destination\_factory\_jndi\_name* JNDI name: *destination\_jndi\_name* Exception message: *exception\_message*

**Explanation**

Events that are to be published to the event group will be rolled back until a connection can be made or until configuration of the event group is modified.

**Administrator Response**

Ensure that the event group definition contains a valid JMS destination.

## CEIES0005E

The event distribution failed to create a JMS message for an event notification being sent to the JMS destination with the specified JNDI name. Event group name: *event\_group\_name* Event: *event\_instance* Factory JNDI name: *destination\_factory\_jndi\_name* JNDI name: *destination\_jndi\_name* Exception message: *exception\_message*

### Explanation

See message.

### Administrator Response

Ensure that the event group definition contains a valid JMS destination.

## CEIES0006E

The event distribution failed to publish the event notification to the JMS destination with the specified JNDI name. Event group name: *event\_group\_name* Event: *event\_instance* JMS Message: *event\_instance* Factory JNDI name: *destination\_factory\_jndi\_name* JNDI name: *destination\_jndi\_name* Exception message: *exception\_message*

### Explanation

If this is the first destination configured for the event group, the event is rolled back to the sender. If this is not the first destination configured for the event group, event distribution attempts to publish the event to the other destinations for this event group.

### Administrator Response

Ensure that the event group definition contains a valid JMS destination.

## CEIES0007W

The event distribution failed to disconnect from the JMS destination with the specified JNDI name. Event group name: *event\_group\_name* Factory JNDI name: *destination\_factory\_jndi\_name* JNDI name: *destination\_jndi\_name* Exception message: *exception\_message*

### Explanation

No recovery is done because the JMS destination is closed.

## CEIES0008W

The event bus failed to store the event in the data store because the global instance ID in the event already exists in the data store. The event was not persisted in the configured data store, and it was not published to any message consumers. Event: *event\_instance* Exception message: *exception\_message*

### Explanation

The global instance ID uniquely identifies the event in the data store. It is assumed that the passed event is a duplicate of an event that already exists in the data store. Therefore, the event is not processed, and the current transaction is not marked for rollback.

### **Administrator Response**

If this is a recurring problem, inspect the source of the event to determine if the source is creating duplicate global instance IDs.

## **CEIES0009E**

The event bus failed to store the event in the data store because the data store could not be found using the specified JNDI name. The event was not persisted in the configured data store, and it was not published to any message consumers. Event: *event\_instance* JNDI name: *data\_store\_jndi\_name* Exception message: *exception\_message*

### **Explanation**

The event is not processed, and the current transaction is marked for rollback.

### **Programmer Response**

Ensure that the data store exists at the specified JNDI name. Also, see the exception message for additional information about the cause of the problem.

## **CEIES0010E**

The event bus failed to store the event in the data store referenced by the JNDI name. The event was not persisted in the configured data store, and it was not published to any message consumers. Event: *event\_instance* JNDI name: *data\_store\_jndi\_name* Exception message: *exception\_message*

### **Explanation**

The event is not processed, and the current transaction is marked for rollback.

### **Administrator Response**

See the exception message for additional information about the cause of the problem.

## **CEIES0011E**

The event bus failed to distribute the event. The event was not persisted in the configured data store, and it was not published to any message consumers. Event: *event\_instance* Exception message: *exception\_message*

### **Explanation**

The event is not processed, and the current transaction is marked for rollback.

### **Administrator Response**

See the exception message for additional information about the cause of the problem.

## **CEIES0012E**

The JMS type could not be found when processing a JMS message. JMS message: *message*

**Explanation**

See message.

**Administrator Response**

If this is a recurring problem, inspect the source of the event to determine if the source is creating JMS messages that are not valid.

**CEIES0013E**

The event could not be found when processing a JMS message. JMS message: *message*

**Explanation**

See message.

**Administrator Response**

If this is a recurring problem, inspect the source of the event to determine if the source is creating JMS messages that are not valid.

**CEIES0014W**

The event was not retrieved because the JMS message type is not valid. JMS message: *message*

**Explanation**

See message.

**Administrator Response**

If this is a recurring problem, inspect the source of the event to determine if the source is creating JMS messages that are not valid.

**CEIES0015E**

Creation of a JMS message for the event failed. Event: *event\_instance* Exception message: *exception\_message*

**Explanation**

See message.

**Administrator Response**

See the exception message for additional information about the cause of the problem.

**CEIES0016E**

The notification helper failed to set the event selector because the event selector is not in the correct format. Event selector: *event\_selector* Exception message: *exception\_message*

**Explanation**

See message.

**Administrator Response**

Ensure that the calling application is passing a valid event selector.

**CEIES0018E**

The event bus failed to process the event because the event server does not support the event version specified in the event. Event: *event\_instance* Event version: *event\_version* Event server version: *event\_server\_version*

**Explanation**

The event is not processed, and the current transaction is marked for rollback.

**Administrator Response**

If this is a recurring problem, inspect the source of the event to determine the cause of the problem.

**CEIES0019E**

The method *method* called by an event consumer on event access failed because the data store component is not enabled.

**Explanation**

See message.

**Administrator Response**

If the data store should be enabled, update the configuration of the Common Event Infrastructure application.

**CEIES0020E**

The method *method* called by an event consumer on event access failed. Exception message: *exception\_message*

**Explanation**

See message.

**Administrator Response**

See the exception message for additional information about the cause of the problem.

**CEIES0021E**

The home interface of the configured data store could not be found using the specified JNDI name. JNDI name: *jndi\_name* Exception message: *exception\_message*



**Explanation**

See message.

**Administrator Response**

Ensure that the correct JNDI name is configured for the data store. Also, ensure that the data store is deployed to the specified JNDI name.

**CEIES0023E**

The event server failed to create an instance of the configured data store because the data store version and the event server version are not compatible. Data store version: *datastore\_cbe\_version* Event server version: *eventserver\_cbe\_version*

**Explanation**

See message.

**Administrator Response**

The event server must use a data store that supports the same event version.

**CEIES0024E**

The event server failed to create an instance of the configured data store because the component metadata could not be retrieved from the data store. Exception message: *exception\_message*

**Explanation**

See message.

**Administrator Response**

See the exception message for additional information about the cause of the problem.

**CEIES0025E**

The event server failed to create an instance of the configured data store. Exception message: *exception\_message*

**Explanation**

See message.

**Administrator Response**

See the exception message for additional information about the cause of the problem.

**CEIES0026E**

The configuration of the event server is not valid because the data store is enabled, but the data store JNDI name is null. It is not valid to enable the data store and to have a null data store JNDI name. Event server JNDI name: *jndi\_name* Event server configuration object: *configuration\_object*

**Explanation**

See message.

**Administrator Response**

Update the configuration to either disable the data store or provide a valid data store JNDI name.

**CEIES0027E**

The configuration object for the event server at the specified JNDI name could not be found. JNDI name: *jndi\_name* Exception message: *exception\_message*

**Explanation**

See message.

**Administrator Response**

Verify that the Common Event Infrastructure server started. Also, ensure that the failing EJB has a correct EJB reference in the deployment descriptor to the event server configuration object.

**Programmer Response****CEIES0028E**

The configuration object for the event server could not be found when accessing configuration of event group list at the specified JNDI name. JNDI name: *jndi\_name* Exception message: *exception\_message*

**Explanation**

See message.

**Administrator Response**

Verify that the Common Event Infrastructure server started. Also, ensure that the failing EJB has a correct EJB reference in the deployment descriptor to the event server configuration object.

**Programmer Response****CEIES0029E**

The reference obtained from the specified JNDI name is not a reference to a notification helper factory. JNDI name: *jndi\_name* Exception message: *exception\_message*

**Explanation**

An application attempted to look up a notification helper factory from a JNDI namespace, but the object found in the specified location was not a reference to a notification helper factory.

**Administrator Response**

See the exception message for additional information about the cause of the problem.

## CEIES0030W

The event bus message driven bean failed to decode the JMS message because the message did not contain a valid message type. JMS message: *message* Message type: *message\_type*

### Explanation

The message is not processed, and the current transaction is marked for rollback causing the message to be returned to the JMS provider for handling.

### Administrator Response

If this is a recurring problem, check the source of the message to determine the cause of the problem.

## CEIES0031E

The event bus message driven bean failed to decode a JMS message. JMS message: *message* Exception message: *exception\_message*

### Explanation

The event is not processed, and the current transaction is marked for rollback.

### Administrator Response

See the exception message for additional information about the cause of the problem.

## CEIES0032E

The event bus message driven bean failed to submit an event to the event bus. Event: *event* Exception message: *exception\_message*

### Explanation

The event is not processed, and the current transaction is marked for rollback.

### Administrator Response

See the exception message for additional information about the cause of the problem.

## CEIES0033E

The event bus message driven bean failed to process a JMS message. JMS message: *message* Exception message: *exception\_message*

### Explanation

The message is not processed, and the current transaction is marked for rollback.

### Administrator Response

See the exception message for additional information about the cause of the problem.

## CEIES0034E

The event bus message driven bean failed because the home interface of the event bus could not be found using the specified JNDI name. JNDI name: *jndi\_name* Exception message: *exception\_message*

### Explanation

See message.

### Administrator Response

Ensure that the event bus message driven bean has a correct EJB reference in the deployment descriptor to the event bus.

## CEIES0035E

The event bus message driven bean failed to create an instance of the event bus. Exception message: *exception\_message*

### Explanation

See message.

### Administrator Response

See the exception message for additional information about the cause of the problem.

## CEIES0036E:

The event bus message driven bean failed to create an instance of the event bus. The event was not persisted, and it was not published to any message consumers. Event: *event\_instance* Exception message: *exception\_message*

### Explanation

The event is not processed, and the current transaction is marked for rollback.

### Administrator Response

See the exception message for additional information about the cause of the problem.

## CEIES0037E

The event selector parser failed to parse the event selector because it did not contain a valid XPath expression. Event selector: *event\_selector* Exception message: *exception\_message*

### Explanation

See message.

### Administrator Response

Ensure that the configuration setting for the event selector contains a valid XPath expression.

## CEIES0038E

The event selector does not represent a valid event selector for filtering events because it did not contain a valid XPath expression. Event selector: *event\_selector*

### Explanation

The event selector used by either event distribution or the notification helper must contain a valid XPath expression, it must select on `CommonBaseEvent`, and it must contain at least one predicate.

### Administrator Response

Ensure that the configuration setting for the event selector contains a valid XPath expression.

## CEIES0039E

The event selector subexpression is not valid because the `creationTime` property must be compared to a valid `xsd:datetime` value. Event selector: *event\_selector*

### Explanation

See message.

### Administrator Response

Ensure that the configuration setting for the event selector contains a valid XPath expression.

## CEIES0040E

The event selector subexpression is not valid because it contains an event selector expression that is not supported. Event selector: *event\_selector*

### Explanation

See message.

### Administrator Response

Ensure that the configuration setting for for the event selector contains a valid XPath expression.

## CEIES0041E

The event selector failed to parse using the XPath parser, and the notification helper failed to set the event selector. Event selector: *event\_selector* Exception message: *exception\_message*

### Explanation

See message.

### Administrator Response

Ensure that the configuration setting for the event selector contains a valid XPath expression.

## **CEIES0042E**

The notification helper failed because the event group configuration could not be found in JNDI. Event group: *event\_group* JNDI name: *jndi\_name* Exception message: *exception\_message*

### **Explanation**

See message.

### **Administrator Response**

Ensure that the event group is configured correctly.

## **CEIES0043E**

The notification helper failed because the event group configuration could not be found using the specified JNDI name. Event group: *event\_group* JNDI name: *jndi\_name* Exception message: *exception\_message*

### **Explanation**

See message.

### **Administrator Response**

Ensure that the event group is configured correctly.

## **CEIES0044E**

The home interface of the event distribution could not be found using the specified JNDI name. JNDI name: *jndi\_name* Exception message: *exception\_message*

### **Explanation**

See message.

### **Administrator Response**

Ensure that the event bus has a correct EJB reference in the deployment descriptor to the event distribution.

## **CEIES0045E**

The event server failed to create an instance of the event distribution. Exception message: *exception\_message*

### **Explanation**

See message.

### **Administrator Response**

See the exception message for additional information about the cause of the problem.

## **CEIES0046E:**

The event distribution failed because the configuration of the event group list could not be found using the specified JNDI name. JNDI name: *jndi\_name* Exception message: *exception\_message*

### **Explanation**

See message.

### **Administrator Response**

Ensure that the correct JNDI name is configured for the event group list.

## **CEIES0048E**

The event group is not defined in the event group list that the event server instance is using. Event group: *event\_group\_name*

### **Explanation**

See message.

### **Administrator Response**

Ensure that the event group is defined for this instance of the event server. Also, restart the application server that is running the instance of the event server to incorporate the latest configuration.

## **CEIES0049E**

The event group does not contain a valid event selector. Event group: *event\_group\_name* Event selector: *event\_selector*

### **Explanation**

See message.

### **Administrator Response**

Correct the specified event selector in the event group definition. Also, restart the application server that is running the instance of the event server to incorporate the latest configuration.

## **CEIES0050E**

The event selector is not valid. Event selector: *event\_selector*

### **Explanation**

See message.

### **Administrator Response**

Correct the event selector in the calling application.



## **CEIIN0001E**

The directory *directory\_name* failed to be created.

### **Explanation**

See message.

### **Administrator Response**

Consult the Common Event Infrastructure log files to determine the cause of the problem.

## **CEIIN0002I**

The enterprise application was successfully deployed to the specified node. Application: *application\_name*  
Node: *node\_name* Server: *server\_name*

### **Explanation**

No additional information is available for this message.

## **CEIIN0003E**

The enterprise application failed to be deployed to the specified node. Application: *application\_name* Node:  
*node\_name* Server: *server\_name*

### **Explanation**

See message.

### **Administrator Response**

Consult the Common Event Infrastructure log files to determine the cause of the problem.

## **CEIIN0004I**

The enterprise application was uninstalled successfully from the specified node. Application:  
*application\_name* Node: *node\_name*

### **Explanation**

No additional information is available for this message.

## **CEIIN0005E**

The enterprise application failed to be uninstalled from the specified node. Application: *application\_name*  
Node: *node\_name*

### **Explanation**

See message.

**Administrator Response**

Consult the Common Event Infrastructure log files to determine the cause of the problem.

**CEIIN0006E**

The file *file\_name* was not found in the specified location.

**Explanation**

See message.

**Administrator Response**

Consult the Common Event Infrastructure log files to determine the cause of the problem.

**CEIIN0007E**

The Common Event Infrastructure database scripts for DB2 failed to be created.

**Explanation**

See message.

**Administrator Response**

Consult the Common Event Infrastructure log files to determine the cause of the problem.

**CEIIN0008E**

Creation of the DB2 database for the Common Event Infrastructure application failed.

**Explanation**

See message.

**Administrator Response**

Consult the Common Event Infrastructure log files to determine the cause of the problem.

**CEIIN0009E**

The DB2 data source for the Common Event Infrastructure database failed to be created.

**Explanation**

See message.

**Administrator Response**

Consult the Common Event Infrastructure log files to determine the cause of the problem.

## **CEIIN0010E**

The Common Event Infrastructure database scripts for Oracle failed to be created.

### **Explanation**

See message.

### **Administrator Response**

Consult the Common Event Infrastructure log files to determine the cause of the problem.

## **CEIIN0011E**

Creation of the Oracle database for the Common Event Infrastructure application failed.

### **Explanation**

See message.

### **Administrator Response**

Consult the Common Event Infrastructure log files to determine the cause of the problem.

## **CEIIN0012E**

The Oracle data source for the Common Event Infrastructure database failed to be created.

### **Explanation**

See message.

### **Administrator Response**

Consult the Common Event Infrastructure log files to determine the cause of the problem.

## **CEIIN0013E**

The Common Event Infrastructure database scripts for CloudScape failed to be created.

### **Explanation**

See message.

### **Administrator Response**

Consult the Common Event Infrastructure log files to determine the cause of the problem.

## **CEIIN0014E**

Creation of the CloudScape database for the Common Event Infrastructure application failed.

**Explanation**

See message.

**Administrator Response**

Consult the Common Event Infrastructure log files to determine the cause of the problem.

**CEIIN0015E**

The CloudScape data source for the Common Event Infrastructure database failed to be created.

**Explanation**

See message.

**Administrator Response**

Consult the Common Event Infrastructure log files to determine the cause of the problem.

**CEIIN0016E**

The response file *file\_name* is not valid because it does not contain the required parameters, or it contains values that are not valid.

**Explanation**

See message.

**Administrator Response**

Consult the Common Event Infrastructure log files to determine the cause of the problem.

**CEIIN0017E**

A required parameter is missing from the response file. Required parameter: *parameter\_name*

**Explanation**

See message.

**Programmer Response**

Ensure that the response file contains the specified parameter.

**CEIIN0018I**

The Common Event Infrastructure configuration files were uninstalled successfully.

**Explanation**

No additional information is available for this message.

## **CEIIN0019W**

The directory *directory\_name* was not found in the specified location.

### **Explanation**

See message.

### **Administrator Response**

Consult the Common Event Infrastructure log files to determine the cause of the problem.

## **CEIIN0020E**

The DB\_TYPE parameter is not specified in the response file *file\_name*.

### **Explanation**

See message.

### **Administrator Response**

Ensure that the DB\_TYPE parameter is specified in the response file.

## **CEIIN0021E**

The DB\_TYPE parameter does not contain a valid value. Valid values are: *validValues*

### **Explanation**

See message.

### **Administrator Response**

Ensure that the DB\_TYPE parameter that is specified in the response file contains a valid value.

## **CEIIN0022E**

The PAGE\_CACHE\_SIZE parameter in the response file contains a value that is not valid.

### **Explanation**

The PAGE\_CACHE\_SIZE parameter specifies the number of memory pages to use to cache data. Set the PAGE\_CACHE\_SIZE parameter to a number greater than 40.

### **Administrator Response**

Ensure that the PAGE\_CACHE\_SIZE parameter contains a numeric value.

## **CEIIN0023E**

The installation of the Common Event Infrastructure application completed, but errors occurred.

**Explanation**

See message.

**Administrator Response**

Consult the Common Event Infrastructure log files to determine the cause of the problem.

**CEIIN0024E**

The parameter *parameter\_name* in the response file contains a value *invalidValue*, which is not valid. Value values are: *validValues*.

**Explanation**

See message.

**Administrator Response**

Change the value of the specified parameter to a valid value.

**CEIIN0025E**

The CloudScape database already exists in the specified directory path. Database name: *dbName*  
Directory path: *path*

**Explanation**

See message.

**Administrator Response**

Remove the database before attempting to create a database with the same name.

**CEIIN0026I**

The enterprise application must be deployed to use the Common Event Infrastructure server. Application: *application\_name*

**Explanation**

See message.

**Administrator Response**

Deploy the specified application manually using the wsadmin tool.

**CEIIN0027E**

Creation of the Common Event Infrastructure database failed.

**Explanation**

See message.

### **Administrator Response**

Consult the Common Event Infrastructure log files to determine the cause of the problem.

### **CEIIN0028E**

When running the Common Event Infrastructure installer from the command line, the correct command syntax was not provided.

#### **Explanation**

See message.

### **Programmer Response**

Provide the correct command syntax.

### **CEIIN0029E**

The response file *file\_name* does not exist.

#### **Explanation**

See message.

### **Administrator Response**

Ensure that you specify the correct response file name and location.

### **CEIIN0030E**

The Common Event Infrastructure server is not installed.

#### **Explanation**

No additional information is available for this message.

### **CEIIN0031E**

The directory *directory\_name* is not a valid WAS\_HOME directory.

#### **Explanation**

A WebSphere Application Server installation is not found in the specified directory.

### **Administrator Response**

Ensure that you have specified the correct location for the WebSphere Application Server installation.

### **CEIIN0032E**

The languages specified in the response file *file\_name* are not a valid or are not supported.



**Explanation**

See message.

**Programmer Response**

Ensure that you specify only supported languages in the response file.

**CEIIN0033E**

The WebSphere Application Server product version files *file\_name1* and *file\_name2* are not found.

**Explanation**

See message.

**Administrator Response**

Verify that the installation of the WebSphere Application Server product completed successfully.

**CEIIN0034E**

The Common Event Infrastructure server is already installed.

**Explanation**

No additional information is available for this message.

**CEIIN0035W**

The enterprise application has already been deployed to the specified node. Application: *application\_name*  
Node: *node\_name* Server: *server\_name*

**Explanation**

See message.

**CEIIN0036W**

The enterprise application has not been deployed to the specified node. Application: *application\_name*  
Node: *node\_name*

**Explanation**

No additional information is available for this message.

**CEIIN0037E**

The file *file\_name* could not be created.

**Explanation**

See message.

**Administrator Response**

Consult the Common Event Infrastructure log files to determine the cause of the problem.

**CEIIN0038E**

The file *file\_name1* could not be copied to file *file\_name2*.

**Explanation**

See message.

**Administrator Response**

Consult the Common Event Infrastructure log files to determine the cause of the problem.

**CEIIN0039E**

The file *file\_name* could not be placed in the directory *directory\_name*.

**Explanation**

See message.

**Administrator Response**

Consult the Common Event Infrastructure log files to determine the cause of the problem.

**CEIIN0040E**

The WebSphere Application Server configuration directory *directory\_name* was not found.

**Explanation**

See message.

**Administrator Response**

Verify that the installation of the WebSphere Application Server product completed successfully.

**CEIIN0041E**

The WebSphere Application Server configuration directory *directory\_name* does not contain any files.

**Explanation**

See message.

**Administrator Response**

Verify that the installation of the WebSphere Application Server product completed successfully.

## **CEIIN0042I**

The installation of the Common Event Infrastructure application completed successfully.

### **Explanation**

No additional information is available for this message.

## **CEIIN0043E**

The installation of the Common Event Infrastructure application failed.

### **Explanation**

See message.

### **Administrator Response**

Consult the Common Event Infrastructure log files to determine the cause of the problem.

## **CEIIN0044E**

The enterprise application failed to uninstall because the uninstallation script was not found. Application: *application\_name* Script name: *script\_name*

### **Explanation**

See message.

### **Administrator Response**

Verify that the installation of the Common Event Infrastructure application completed successfully.

## **CEIIN0045I**

The installation of the Common Event Infrastructure configuration files completed successfully.

### **Explanation**

No additional information is available for this message.

## **CEIIN0046E**

Installation of the Common Event Infrastructure configuration files failed.

### **Explanation**

See message.

### **Administrator Response**

Consult the Common Event Infrastructure log files to determine the cause of the problem.

## **CEIIN0047E**

Uninstallation of the Common Event Infrastructure configuration files failed.

### **Explanation**

See message.

### **Administrator Response**

Consult the Common Event Infrastructure log files to determine the cause of the problem.

## **CEIIN0048E**

The WebSphere Application Server *server\_name* is stopped. Start the server to configure the Common Event Infrastructure JDBC provider and data source.

### **Explanation**

See message.

### **Administrator Response**

See message.

## **CEIIN0049I**

Version *version\_number* of the Common Event Infrastructure application is installed.

### **Explanation**

No additional information is available for this message.

## **CEIIN0050W**

Version *version\_number* of the Common Event Infrastructure application is installed, but this version is not valid.

### **Explanation**

See message.

### **Administrator Response**

Verify that the installation of the Common Event Infrastructure application completed successfully.

## **CEIIN0051I**

Installing version *version\_number* of the Common Event Infrastructure application.

### **Explanation**

No additional information is available for this message.

## **CEIIN0052I**

Uninstalling version *version\_number* of the Common Event Infrastructure application.

### **Explanation**

No additional information is available for this message.

## **CEIIN0053I**

The installation of the Common Event Infrastructure library files completed successfully.

### **Explanation**

No additional information is available for this message.

## **CEIIN0054E**

The installation of the Common Event Infrastructure library files failed.

### **Explanation**

See message.

### **Administrator Response**

Consult the Common Event Infrastructure log files to determine the cause of the problem.

## **CEIIN0055I**

Uninstallation of the Common Event Infrastructure library files completed successfully.

### **Explanation**

No additional information is available for this message.

## **CEIIN0056E**

Uninstallation of the Common Event Infrastructure library files failed.

### **Explanation**

See message.

### **Administrator Response**

Consult the event infrastructure log files to determine the cause of the problem.

## **CEIIN0057I**

Installation of the Common Event Infrastructure tool files completed successfully.

**Explanation**

No additional information is available for this message.

**CEIIN0058I**

Installation of the Common Event Infrastructure tool files failed.

**Explanation**

See message.

**Administrator Response**

Consult the Common Event Infrastructure log files to determine the cause of the problem.

**CEIIN0059I**

Uninstallation of the Common Event Infrastructure tool files completed successfully.

**Explanation**

No additional information is available for this message.

**CEIIN0060E**

Uninstallation of the Common Event Infrastructure tool files failed.

**Explanation**

See message.

**Administrator Response**

Consult the Common Event Infrastructure log files to determine the cause of the problem.

**CEIIN0061W**

Setting execute permissions on some files in the directory *directory\_name* failed. Manually set the execute permissions for executable scripts before running the scripts.

**Explanation**

See message.

**Administrator Response**

See message.

**CEIIN0062I**

Uninstallation of the Common Event Infrastructure application completed successfully.

**Explanation**

No additional information is available for this message.

**CEIIN0063E**

Uninstallation of the Common Event Infrastructure application failed.

**Explanation**

See message.

**Administrator Response**

Consult the Common Event Infrastructure log files to determine the cause of the problem.

**CEIIN0064E**

The Common Event Infrastructure version file *file\_name* could not be deleted. Manually delete the specified file before reinstalling the Common Event Infrastructure application.

**Explanation**

See message.

**Administrator Response**

See message.

**CEIIN0065I**

Uninstallation of the common event infrastructure database completed successfully.

**Explanation**

No additional information is available for this message.

**CEIIN0066E**

Uninstallation of the Common Event Infrastructure database failed.

**Explanation**

See message.

**Administrator Response**

Consult the Common Event Infrastructure log files to determine the cause of the problem.

**CEIIN0067E**

The WebSphere Network Deployment Manager is not running. The WebSphere Network Deployment Manager must be running before configuring or uninstalling the Common Event Infrastructure application.



**Explanation**

See message.

**Administrator Response**

Start the WebSphere Network Deployment Manager.

**CEIIN0068E**

An error occurred when installing the Common Event Infrastructure application.

**Explanation**

See message.

**Administrator Response**

Consult the Common Event Infrastructure log files to determine the cause of the problem.

**CEIIN0069E**

An error occurred when uninstalling the Common Event Infrastructure application.

**Explanation**

See message.

**Administrator Response**

Consult the Common Event Infrastructure log files to determine the cause of the problem.

**CEIIN0070I**

Consult the log file *log\_file\_name* for more information.

**Explanation**

No additional information is available for this message.

**CEIIN0071I**

The installation of the Common Event Infrastructure administrative panel completed successfully.

**Explanation**

No additional information is available for this message.

**CEIIN0072E**

The installation of the Common Event Infrastructure administrative panel failed.

**Explanation**

See message.

**Administrator Response**

Consult the Common Event Infrastructure log files to determine the cause of the problem.

**CEIIN0073I**

The uninstallation of the Common Event Infrastructure administrative panel completed successfully.

**Explanation**

No additional information is available for this message.

**CEIIN0074E**

The uninstallation of the Common Event Infrastructure administrative panel failed.

**Explanation**

See message.

**Administrator Response**

Consult the Common Event Infrastructure log files to determine the cause of the problem.

**CEIIN0075I**

The Common Event Infrastructure administrative panel is not installed.

**Explanation**

No additional information is available for this message.

**CEIIN0076E**

The file *file\_name* could not be deleted.

**Explanation**

See message.

**Administrator Response**

Manually delete the specified file.

**CEIIN0077E**

The directory *directory* could not be deleted.

**Explanation**

See message.

**Administrator Response**

Manually delete the specified directory.

**CEIIN0078E**

Writing a file in EBCDIC format failed.

**Explanation**

See message.

**Administrator Response**

Consult the Common Event Infrastructure log files to determine the cause of the problem.

**CEIIN0079E**

Reading a file in EBCDIC format failed.

**Explanation**

See message.

**Administrator Response**

Consult the Common Event Infrastructure log files to determine the cause of the problem.

**CEIIN0083E**

The system could not determine if the WebSphere Application Server is operating in a standalone or a network environment.

**Explanation**

See message.

**Administrator Response**

Consult the Common Event Infrastructure log files to determine the cause of the problem.

**CEIIN0084E**

The enterprise application, *application\_name*, could not be uninstalled because WebSphere Application Server security is enabled.

**Explanation**

See message.

### **Administrator Response**

Manually uninstall the specified application.

### **CEIIN0085I**

The Common Event Infrastructure enterprise application could not be uninstalled. Manually uninstall the enterprise application *application\_name*, and then delete the directory *directory\_name*.

### **Explanation**

See message.

### **Administrator Response**

Consult the event infrastructure log files to determine the cause of the problem and manually uninstall the specified enterprise application.

### **CEIIN0086E**

The Common Event Infrastructure data source could not be uninstalled because WebSphere Application Server security is enabled.

### **Explanation**

See message.

### **Administrator Response**

Manually uninstall the Common Event Infrastructure data source.

### **CEIIN0087I**

The Common Event Infrastructure database or data source could not be uninstalled. Manually uninstall the Common Event Infrastructure database and data source. Then delete the directory *directory\_name*.

### **Explanation**

See message.

### **Administrator Response**

Consult the Common Event Infrastructure log files to determine the cause of the problem. Manually uninstall the Common Event Infrastructure database and data source.

### **CEIIN0088E**

The LOG\_FILE\_SIZE parameter in the response file contains the value *value*, which is not valid.

### **Explanation**

The LOG\_FILE\_SIZE parameter defines the size of each primary and secondary log file in 4K pages. The log file size limits the number of log records that can be written to the log file before it becomes full and a new log file is required.

### **Administrator Response**

Enter a valid value in the response file for the LOG\_FILE\_SIZE parameter. Refer to the appropriate DB2 documentation for additional information.

## **CEIIN0089E**

The NUM\_PRIMARY\_LOG parameter in the response file contains the value *value*, which is not valid.

### **Explanation**

The primary log files establish a fixed amount of storage allocated to the recovery log files. Use the NUM\_PRIMARY\_LOG parameter to specify the number of primary log files to be preallocated.

### **Administrator Response**

Enter a valid value in the response file for the NUM\_PRIMARY\_LOG parameter. Refer to the appropriate DB2 documentation for additional information.

## **CEIIN0090E**

The NUM\_SECONDARY\_LOG parameter in the response file contains the value *value*, which is not valid.

### **Explanation**

The NUM\_SECONDARY\_LOG parameter specifies the number of secondary log files that are created and used for recovery log files, if needed.

### **Administrator Response**

Enter a valid value in the response file for the NUM\_SECONDARY\_LOG parameter in the response file. Refer to the appropriate DB2 documentation for more information.

## **CEIIN0091E**

An error occurred when running the PluginProcessor command.

### **Explanation**

See message.

### **Administrator Response**

Perform the following steps:

1. Manually delete the Common Event Infrastructure files that were not deleted.
2. Run the following command: **\$WAS\_HOME/bin/PluginProcessor -restore**

## **CEIIN0092I**

You must manually restore the Administrative console enterprise application.

### **Explanation**

See message.

## Administrator Response

Perform the following steps:

1. Delete the Common Event Infrastructure files that were not deleted.
2. Run the following command: **\$WAS\_HOME/bin/PluginProcessor -restore**

## CEIIN0093I

To restore the system, copy the file *source\_file\_name* manually to the file *destination\_file\_name*.

### Explanation

No additional information is available for this message.

## CEIIN0094I

To restore the system, manually collapse the enterprise application directory *directory\_name* to the temporary file *source\_file\_name* using the `$WAS_HOME/bin/EARExpander` command. Replace the file *destination\_file\_name* with the temporary file.

### Explanation

No additional information is available for this message.

## CEIIN0095E

The backup of the file *source\_file\_name* to the file *destination\_file\_name* failed.

### Explanation

See message.

## Administrator Response

Consult the Common Event Infrastructure log files to determine the cause of the problem.

## CEIIN0096E

The file *file\_name* cannot be restored from the backup copy of the file *backup\_file\_name*.

### Explanation

See message.

## Administrator Response

Consult the Common Event Infrastructure log files to determine the cause of the problem. Manually restore the file from the backup copy of the file.

## CEIIN0097E

Creation of a collapsed format of the Administrative console enterprise application in the directory *directory\_name* to the file *file\_name* failed.

**Explanation**

See message.

**Administrator Response**

Consult the Common Event Infrastructure log files to determine the cause of the problem.

**CEIIN0098E**

The database configuration script file, *file\_name*, was not found.

**Explanation**

See message.

**Administrator Response**

Verify that the preceding steps in the installation of the Common Event Infrastructure application completed successfully.

**CEIIN0099I**

The installation of the Common Event Infrastructure database completed successfully.

**Explanation**

No additional information is available for this message.

**CEIIN0100E**

The installation of the Common Event Infrastructure database failed.

**Explanation**

See message.

**Administrator Response**

Consult the Common Event Infrastructure log files to determine the cause of the problem. Manually create the Common Event Infrastructure database.

**CEIIN0101I**

Creation of the default Common Event Infrastructure database during installation is supported only in a WebSphere Application Server standalone environment.

**Explanation**

See message.



## **Administrator Response**

### **CEIIN0102I**

You must manually create the Common Event Infrastructure database using the database configuration scripts.

#### **Explanation**

No additional information is available for this message.

### **CEIIN0103I**

Deployment of the Common Event Infrastructure enterprise application during installation is supported only in a WebSphere Application Server standalone environment. Application: *application\_name*

#### **Explanation**

No additional information is available for this message.

### **CEIIN0104E**

Configuration of the default Common Event Infrastructure service failed.

#### **Explanation**

See message.

## **Administrator Response**

Consult the Common Event Infrastructure log files to determine the cause of the problem. Manually configure the Common Event Infrastructure service.

---

## Chapter 24. Using the transaction service

These topics provide information about using transactions with WebSphere applications

WebSphere applications can use transactions to coordinate multiple updates to resources as atomic units (as indivisible units of work) such that all or none of the updates are made permanent.

In WebSphere Application Server, transactions are handled by three main components:

- A transaction manager that supports the enlistment of recoverable XAResources and ensures that each such resource is driven to a consistent outcome either at the end of a transaction or after a failure and restart of the application server.
- A container in which the J2EE application runs. The container manages the enlistment of XAResources on behalf of the application when the application performs updates to transactional resource managers (for example, databases). Optionally, the container can control the demarcation of transactions for enterprise beans configured for container-managed transactions.
- An application programming interface (UserTransaction) that is available to bean-managed enterprise beans and servlets. This allows such application components to control the demarcation of their own transactions.

For more information about using transactions with WebSphere applications, see the following topics:

- Transaction support in WebSphere Application Server
- Using local transactions
- Developing a WebSphere application to use transactions
- Configuring transaction properties for an application server
- Managing active transactions
- Managing transaction logging for optimum server availability
- Interoperating transactionally between application servers
- Troubleshooting transactions
- Transaction service exceptions
- UserTransaction interface - methods available
- Coordinating access to 1-PC and 2-PC-capable resources within the same transaction
- Implementing WebSphere enterprise applications that use ActivitySessions

---

### Transaction support in WebSphere Application Server

A transaction is unit of activity within which multiple updates to resources can be made atomic (as an indivisible unit of work) such that all or none of the updates are made permanent. For example, multiple SQL statements to a relational database are committed atomically by the database during the processing of an SQL COMMIT statement. In this case, the transaction is contained entirely within the database manager and can be thought of as a resource manager local transaction (RMLT). In some contexts, a transaction is referred to as a logical unit of work (LUW).

The way that applications use transactions depends on the type of application component, as follows:

- A session bean can either use container-managed transactions (where the bean delegates management of transactions to the container) or bean-managed transactions (where the bean manages transactions itself).
- Entity beans use container-managed transactions.
- Web components (servlets) use bean-managed transactions.

WebSphere Application Server is a transaction manager that supports the coordination of resource managers through their XAResource interface and participates in distributed global transactions with other OTS 1.2 compliant transaction managers (for example J2EE 1.3 application servers). WebSphere applications can also be configured to interact with databases, JMS queues, and JCA connectors through their local transaction support when distributed transaction coordination is not required.

Resource managers that offer transaction support can be categorized into those that support two-phase coordination (by offering an XAResource interface) and those that support only one-phase coordination (for example through a LocalTransaction interface). The WebSphere Application Server transaction support provides coordination, within a transaction, for any number of two-phase capable resource managers. It also enables a single one-phase capable resource manager to be used within a transaction in the absence of any other resource managers, although a WebSphere transaction is not necessary in this case.

Under normal circumstances you cannot mix one-phase commit capable resources and two-phase commit capable resources in the same global transaction, because one-phase commit resources cannot support the prepare phase of two-phase commit. There are some special circumstances where it is possible to include mixed-capability resources in the same global transaction:

- In scenarios where there is only a single one-phase commit resource provider that participates in the transaction and where all the two-phase commit resource-providers that participate in the transaction are used in a read-only fashion. In this case, the two-phase commit resources all vote read-only during the prepare phase of two-phase commit. Because the one-phase commit resource provider is the only provider to actually perform any updates, the one-phase commit resource does not need to be prepared.
- In scenarios where there is only a single one-phase commit resource provider that participates in the transaction with one of more two-phase commit resource providers and where last participant support is available.

**5.1 +** Last participant support (of WBI Server Foundation) enables the use of a single one-phase commit capable resource with any number of two-phase commit capable resources in the same global transaction.

**5.1 +** The ActivitySession service (of WBI Server Foundation) provides an alternative unit-of-work (UOW) scope to that provided by global transaction contexts. It is a distributed context that can be used to coordinate multiple one-phase resource managers. The WebSphere EJB container and deployment tooling support ActivitySessions as an extension to the J2EE programming model. EJBs can be deployed with lifecycles that are influenced by ActivitySession context, as an alternative to transaction context. An application can then interact with a resource manager through its LocalTransaction interface for the period of a client-scoped ActivitySession rather than just the duration of an EJB method.

## Resource manager local transaction (RMLT)

A resource manager local transaction (RMLT) is a resource manager's view of a local transaction; that is, it represents a unit of recovery on a single connection that is managed by the resource manager.

Resource managers include:

- Enterprise Information Systems that are accessed through a resource adapter, as described in the J2EE Connector Architecture 1.0.
- Relational databases that are accessed through a JDBC datasource.
- JMS queue and topic destinations.

Resource managers offer specific interfaces to enable control of their RMLTs. J2EE connector resource adapters that include support for local transactions provide a LocalTransaction interface to enable applications to request that the resource adapter commit or rollback RMLTs. JDBC datasources provide a Connection interface for the same purpose.

The boundary at which all RMLTs must be complete is defined in WebSphere Application Server by a local transaction containment (LTC).

## Global transactions

If an application uses two or more resources, then an external transaction manager is needed to coordinate the updates to both resource managers in a global transaction.

Global transaction support is available to web and enterprise bean J2EE components. Enterprise bean components can be subdivided into beans that exploit container-managed transactions (CMT) or bean-managed transactions (BMT).

BMT enterprise beans and web components can use the Java Transaction API (JTA) UserTransaction interface to define the demarcation of a global transaction. The UserTransaction interface is obtained by a JNDI lookup of `java:comp/UserTransaction`. The UserTransaction is not available to the following components:

- CMT enterprise beans. Any attempt by such beans to obtain the interface results in an exception in accordance with the EJB specification.
- Client applications running outside the Web and EJB containers.

Ensure that programs that perform a JNDI lookup of the UserTransaction interface, use an InitialContext that resolves to a local implementation of the interface. Also ensure that such programs use a JNDI location appropriate for the EJB version.

Before the EJB 1.1 specification, the JNDI location of the UserTransaction interface was not specified. Each EJB container implementor defined it in an implementation-specific manner. Earlier versions of WebSphere Application Server, up to and including Version 3.5.x (without EJB 1.1), bind the UserTransaction interface to a JNDI location of `jta/usertransaction`. WebSphere Application Server Version 4, and later releases, bind the UserTransaction interface at the location defined by EJB 1.1, which is `java:comp/UserTransaction`. WebSphere Application Server, Version 5 no longer provides the `jta/usertransaction` binding within Web and EJB containers to applications at a J2EE level of 1.3 or later. For example, EJB 2.0 applications can use only the `java:comp/UserTransaction` location.

A web component or enterprise bean (CMT or BMT) can get the ExtendedJTATransaction interface through a lookup of `java:comp/websphere/ExtendedJTATransaction`. This interface provides access to the transaction identity and a mechanism to receive notification of transaction completion.

## Local transaction containment (LTC)

A local transaction containment (LTC) is used to define the application server behavior in an unspecified transaction context.

(Unspecified transaction context is defined in the Enterprise JavaBeans 2.0 Specification.)

A LTC is a bounded unit-of-work scope within which zero, one, or more resource manager local transactions (RMLTs) can be accessed. The LTC defines the boundary at which all RMLTs must be complete; any incomplete RMLTs are resolved, according to policy, by the container. An LTC is local to a bean instance; it is not shared across beans even if those beans are managed by the same container. LTCs are started by the container before dispatching a method on a J2EE component (such as an enterprise bean or servlet) whenever the dispatch occurs in the absence of a global transaction context. LTCs are completed by the container depending on the application-configured LTC boundary; for example at the end of the method dispatch. There is no programmatic interface to the LTC support; rather LTCs are managed exclusively by the container and configured by the application deployer through transaction attributes in the application deployment descriptor.

A local transaction containment cannot exist concurrently with a global transaction. If application component dispatch occurs in the absence of a global transaction, the container always establishes an LTC for J2EE components at J2EE 1.3 or later. The only exceptions to this are as follows:

- Where application component dispatch occurs without container interposition; for example, for a stateless session bean create or a servlet-initiated thread.
- J2EE 1.2 web components.
- J2EE 1.2 BMT enterprise beans.

A local transaction containment can be scoped to an `ActivitySession` context that lives longer than the enterprise bean method in which it is started, as described in `ActivitySessions` and transaction contexts.

## Local and global transaction considerations

Applications use resources, such as JDBC data sources or connection factories, that are configured through the Resources view of the WebSphere Application Server Administrative Console. How these resources participate in a global transaction depends on the underlying transaction support of the resource provider. For example, most JDBC providers can provide either XA or non-XA versions of a data source. A non-XA data source can support only resource manager local transactions (RMLTs), but an XA data source can support two-phase commit coordination, as well as local transactions.

If an application uses two or more resource providers that support only RMLTs, then atomicity cannot be assured because of the one-phase nature of these resources. To ensure atomic behavior, the application should use resources that support XA coordination and should access them within a global transaction.

If an application uses only one RMLT, the atomic behavior can be guaranteed by the resource manager, which can be accessed under a local transaction containment context.

An application can also access a single resource manager under a global transaction context, even if that resource manager does not support the XA coordination. An application can do this, because WebSphere Application Server performs an “only resource optimization” and interacts with the resource manager under a RMLT. Within a global transaction context, any attempt to use more than one resource provider that supports only RMLTs causes the global transaction to be rolled back.

At any moment, an instance of an enterprise bean can have work outstanding in either a global transaction context or a local transaction containment context, but never both. An instance of an enterprise bean can change from running under one type of context to the other (in either direction), if all outstanding work in the original context is complete. Any violation of this principle causes an exception to be thrown when the enterprise bean tries to start the new context.

## Extended JTA support

Extended JTA support provides application programming interfaces additional to the `UserTransaction` interface that is defined in the JTA as part of the J2EE specification. Specifically, the API extensions provide the following functionality:

- Access to global and local transaction identifiers associated with the thread.

The global id is based on the `tid` in `CosTransactions::PropagationContext`; and the local id identifies the transaction uniquely within the local JVM.

- A transaction synchronization callback that enables any J2EE component to register an interest in transaction completion.

This can be used by advanced applications to flush updates before transaction completion and clear up state after transaction completion. J2EE (and related) specifications position this function generally as the domain of the J2EE containers.

**5.1 +** The exception is for CMT session beans, for which the EJB specification defines the `SessionSynchronization` interface. WebSphere provides this WebSphere Application Server EnterpriseWBI Server Foundation functionality in recognition that more advanced applications can profit from the ability to receive such notifications.

An application uses a JNDI lookup of `java:comp/websphere/ExtendedJTATransaction` to get an `ExtendedJTATransaction` object, which it then uses as follows:

```
ExtendedJTATransaction exJTA = (ExtendedJTATransaction)ctx.lookup("
 java:comp/websphere/ExtendedJTATransaction");
SynchronizationCallback sync = new SynchronizationCallback();
exJTA.registerSynchronizationCallback(sync);
```

The `ExtendedJTATransaction` object supports the registration of one or more application-provided `SynchronizationCallbacks`. Each registered callback is called at the end of every transaction that runs on the application server (whether the transaction is started locally or imported).

The following information provides an overview of the interfaces provided by the Extended JTA support. For more detailed information, see the Javadoc provided with WBI Server Foundation.

### **SynchronizationCallback interface**

An object implementing this interface is enlisted once through the `ExtendedJTATransaction` interface, and receives notification of the completion of each subsequent transaction mediated by the transaction manager in the local JVM.

Although an object implementing this interface can run in a J2EE server, there is no specific J2EE component active when this object is called. So, the object has limited direct access to any J2EE resources. Specifically, it has no access to the `java:` namespace or to any container-mediated resource. Such an object can cache a reference to a J2EE component (for example, a stateless session bean) that it delegates to. The object would then have all the normal access to J2EE resources and could be used, for example, to acquire a JDBC connection and flush updates to a database during `beforeCompletion`.

### **ExtendedJTATransaction interface**

A WebSphere programming model extension to the J2EE JTA support. An object implementing this interface is bound, by WebSphere J2EE containers that support this interface, at `java:comp/websphere/ExtendedJTATransaction`. Access to this object, when called from an EJB container, is not restricted to bean-managed transactions.

---

## **Developing components to use transactions**

These topics provide information about developing WebSphere application components to use transactions

The way that applications use transactions depends on the type of application component, as follows:

- A session bean can either use container-managed transactions (where the bean delegates management of transactions to the container) or bean-managed transactions (where the bean manages transactions itself).
- Entity beans use container-managed transactions.
- Web components (servlets) use bean-managed transactions.

You configure whether a component uses container- or bean-managed transactions by setting an appropriate value on the `Transaction type` deployment attribute, as described in [Configuring transactional deployment attributes using the Assembly Toolkit](#). You can also configure other transactional deployment descriptor attributes.

If you want a session bean to manage its own transactions, you must write the code that explicitly demarcates the boundaries of a transaction as described in [Using bean-managed transactions](#).

Similarly, if you want a Web component to use transactions, you must write the code that explicitly demarcates the boundaries of a transaction as described in [Using bean-managed transactions](#).

## **Configuring transactional deployment attributes using the Assembly Toolkit**

Use this task to configure the transactional deployment descriptor attributes associated with an EJB or Web module, to enable a J2EE application to use transactions.



This topic describes the use of the Assembly Toolkit to configure the deployment attributes of an application. This task description assumes that you have an EAR file for an application component, that can be deployed in WebSphere Application Server. For more details about using the Assembly Toolkit, see *Assembling applications with the Assembly Toolkit*.

To set transactional attributes in the deployment descriptor for an application component (enterprise bean or servlet), complete the following steps:

1. Start the Assembly Toolkit.
2. Create or edit the application EAR file. For example, to change attributes of an existing application, use the import wizard to import the EAR file into the Assembly Toolkit. To start the import wizard:
  - a. Click **File-> Import-> EAR file**
  - b. Click **Next**, then select the EAR file.
  - c. Click **Finish**.
3. In the J2EE Hierarchy view, right-click the component instance, then click **Open With > Deployment Descriptor Editor**. For example:
  - For a session bean, expand **EJB Modules-> ejb\_module\_instance-> Session Beans** then select the bean instance.
  - For a servlet, expand **Web Modules-> web\_application-> web\_component** then select the servlet instance.

A property dialog notebook for the component is displayed in the property pane.

4. Set the **Transaction type** attribute, which defines the transactional manner in which the container invokes a method. You can set this attribute to Container or Bean, as follows:
  - For a session bean to use container-managed transactions, set Container
  - For a session bean to use bean-managed transactions, set Bean
  - For an entity bean, set Container
  - For a Web component (servlet), set Bean
5. In the property pane, select the IBM Extensions tab.
6. Under **WebSphere Extensions**, configure J2EE component extensions attributes for extended local transaction containment. To enable management of local transaction containments, configure the following EJB extensions attributes. These attributes configure, for the component, the behavior of the container's local transaction containment (LTC) environment that the container establishes whenever a global transaction is not present.

#### **Boundary**

Specifies the duration of a local transaction context. You can set this attribute to **Bean method** or **ActivitySession**.

**Note:** The ActivitySession option is not supported in the web container.

#### **Resolver**

Specifies how the local transaction is to be resolved before the local transaction context ends: by the application through user code or by the EJB container. You can set this attribute to either **Application** or **ContainerAtBoundary**.

#### **Unresolved action**

Specifies the action that the container must take when the local transaction context scope ends, if resources are uncommitted by an application in a local transaction and the **Resolution control** is set to Application. You can set this attribute to either **Commit** or **Rollback**.

For a value of **Commit**, the container will take the commit action only in the absence of an un-handled exception. If the application method executing under the local transaction context ends with an exception, then the local transaction context is rolled back by the container. (This is the same behavior as for global transactions.)



7. [For EJB components only] For container-managed transactions, configure how the container must manage the transaction boundaries when delegating a method invocation to an enterprise bean's business method:
  - a. In the navigation pane, click the Assembly Descriptor tab. The **Container Transactions** box displays a table of the methods for enterprise beans.
  - b. For each method of the enterprise bean set the **Transaction attribute** attribute to an appropriate value.

If the application uses ActivitySessions, how the container manages transaction boundaries when delegating a method invocation depends on both the **Transaction attribute** attribute, set here, and the **ActivitySession kind** attribute, as described in Configuring ActivitySession deployment attributes for an enterprise JavaBean. For more detail about the relationship between these two properties, see Combining transaction and ActivitySession container policies.
8. Save your changes to the deployment descriptor.
  - a. Close the deployment descriptor editor.
  - b. When prompted, click **Yes** to indicate that you want to save changes to the deployment descriptor.
9. Verify the archive files.
10. Generate code for deployment for EJB modules or for enterprise applications that use EJB modules.
11. Optional: Test your completed module on a WebSphere Application Server installation. Right-click a module, click **Run on Server**, and follow the instructions in the displayed wizard. Note that **Run on Server** works on the Windows, Linux/Intel, and AIX operating systems only; you cannot deploy remotely from the Assembly Toolkit to a WebSphere Application Server installation on a UNIX operating system such as Solaris.

### Important

**Important:** Use **Run On Server** for unit testing only. Assembly Server Toolkit controls the WebSphere Application Server installation and, when an application is published remotely, the Toolkit overwrites the server configuration file for that server. Do not use on production servers.

For instructions on remote testing, see the article "Setting Up a Remote WebSphere Application Server in WebSphere Studio V5" at

[http://www7b.boulder.ibm.com/wsdd/techjournal/0303\\_yuen/yuen.html](http://www7b.boulder.ibm.com/wsdd/techjournal/0303_yuen/yuen.html).

After assembling your application, use a systems management tool to deploy the EAR file onto the application server that is to run the application; for example, using the administrative console as described in Deploying and managing applications.

## Using bean-managed transactions

This topic describes how to enable a session bean or servlet to use bean-managed transactions, to manage its own transactions directly instead of letting the container manage the transactions.

**Note:** Entity beans cannot manage transactions (so cannot use bean-managed transactions).

To enable a session bean or servlet to use bean-managed transactions, complete the following steps:

1. Set the **Transaction type** attribute in the component's deployment descriptor to Bean, as described in Setting transactional attributes in the deployment descriptor.
2. Write the component code to actively manage transactions. When writing the code required by a component to manage its own transactions, remember the following basic rules:
  - An instance of a stateless session bean cannot reuse the same transaction context across multiple methods called by an EJB client.
  - An instance of a stateful session bean can reuse the same transaction context across multiple methods called by an EJB client.

The following code extract shows the standard code required to obtain an object encapsulating the transaction context. There are three basic steps involved:

- The component class must set the value of the `javax.ejb.SessionContext` object reference in the `setSessionContext` method.
- A `javax.transaction.UserTransaction` object is created by calling a lookup on `"java:comp/UserTransaction"`.
- The `UserTransaction` object is used to participate in the transaction by calling transaction methods such as `begin` and `commit` as needed. If an enterprise bean begins a transaction, it must also complete that transaction either by invoking the `commit` method or the `rollback` method.

#### Code example: Getting an object that encapsulates a transaction context

```
...
import javax.transaction.*;
...
public class MyStatelessSessionBean implements SessionBean {
private SessionContext mySessionCtx =null;
...
public void setSessionContext (SessionContext ctx)throws EJBException {
mySessionCtx =ctx;
}
...
public float doSomething(long arg1)throws FinderException,EJBException {
 UserTransaction userTran = (UserTransaction)initCtx.lookup(
 "java:comp/UserTransaction");
 ...
 //User userTran object to call transaction methods
 userTran.begin ();
 //Do transactional work
 ...
 userTran.commit ();
 ...
}
...
}
```

---

## Configuring transaction properties for an application server

Use this task to configure the transaction properties for an application server; for example, to define the location of the directory that contains the transaction log or to change default timeouts associated with transactions.

To configure the transaction properties for an application server, complete the following steps:

1. Start the Administrative console
2. In the navigation pane, select **Servers-> Manage Application Servers-> *your\_app\_server*** This displays the properties of the application server, *your\_app\_server*, in the content pane.
3. Select the Transaction Service tab, to display the properties page for the transaction service, as two notebook pages:

#### Configuration

The values of properties defined in the configuration file. If you change these properties, the new values are applied when the application server next starts.

#### Runtime

The runtime values of properties. If you change these properties, the new values are applied immediately, but are overwritten with the Configuration values when the application server next starts.

4. Select the Configuration tab, to display the transaction-related configuration properties.
5. If you want to change the directory in which transaction logs are written, type the full pathname of the directory in the **Transaction log directory** field. You can check the current runtime value of **Transaction log directory**, by clicking the Runtime tab.

You can also specify a size for the transaction logs, as described in the following step.

**Note:** If you change the transaction log directory, you should apply the change and restart the application server as soon as possible, to minimize the risk of problems caused that might occur before the application server is restarted. For example, if a problem causes the server to fail (with in-flight transactions), the server next starts with the new log directory and is unable to automatically resolve in-flight transactions that were recorded in the old log directory.

6. If you want to change the default file size of transaction log files, modify the **Transaction log directory** field to include a file size setting, in the following format:

*directory\_name;file\_size*

Where

- *directory\_name* is the name of the transaction log directory
- *file\_size* is the new default size specified in bytes. The *nK* or *nM* suffix can be used to indicate kilobytes or megabytes. If you do not specify a file size value, the default value of 1M is used.

For example, `c:\tranlogs;2M` indicates the files are to be created with 2M bytes size and stored in the directory `c:\tranlogs`.

In a non-production environment, you can use the transaction log directory value of `;0` to disable transaction logging. (There must be no directory name element before the size element of 0.) You should not disable transaction logging in a production environment, because this prevents recovery after a system failure and, therefore, data integrity cannot be guaranteed.

7. In the **Total transaction lifetime timeout** field, type the number of seconds a transaction can remain inactive before it is ended by the transaction service. A value of 0 (zero) indicates that there is no timeout limit.
8. In the **Client inactivity timeout** field, type the number of seconds after which a client is considered inactive and the transaction service ends any transactions associated with that client. A value of 0 (zero) indicates that there is no timeout limit.
9. Click **OK**.
10. Stop then restart the application server.

If you change the transaction log directory configuration property to an incorrect directory name, the application server will restart but be unable to open the transaction logs. You should change the configuration property to a valid directory name, then restart the application server.

## Transaction service settings

Use this page to modify transaction service settings.

To view this administrative console page, click **Servers > Application Servers > server > Transaction Service**.

### Transaction log directory

Specifies the name of a directory for this server where the transaction service stores log files for recovery.

A blank value in the server configuration is expanded by the transaction log at startup as the directory `(install_root)/tranlog/(server_name)`.

When the application running on the WebSphere product accesses more than one resource, the WebSphere product stores transaction information to properly coordinate and manage the distributed transaction. In a higher transaction load, this persistence slows down performance of the application server due to its dependency on the operating system and the underlying storage systems.

To achieve better performance, move the transaction log files to a storage device with more physical disk drives, or preferably RAID disk drives. When the log files are moved to the file systems on the RAIDed disks, the task of writing data to the physical media is shared across the multiple disk drives. This allows more concurrent access to persist transaction information and faster access to that data from the logs. Depending upon the design of the application and storage subsystem, performance gains can range from 10% to 100%, or even more in some cases.

This change is applicable only to the configuration where the application uses distributed resources or XA transactions, for example, multiple databases and resources are accessed within a single transaction. Consider setting this property when the application server shows one or more of following signs:

- CPU utilization remains low despite an increase in transactions
- Transactions fail with several time outs
- Transaction rollbacks occur with *unable to enlist transaction* exception
- Application server hangs in middle of a run and requires the server to be restarted
- The disk on which an application server is running shows higher utilization

<b>Data type</b>	String
<b>Default</b>	Initial value is the <code>%WAS_HOME%/tranlog/(server_name)</code> directory and a default size of 1MB.
<b>Recommended</b>	Create a file system with at least 3-4 disk drives RAIDed together in a RAID-0 configuration. Then, create the transaction log on this file system with the default size. When the server is running under load, check the disk input and output. If disk input and output time is more than 5%, consider adding more physical disks to lower the value. If disk input and output is low, but the server is still high, consider increasing the size of the log files.

### Total transaction lifetime timeout

Specifies the maximum duration, in seconds, for transactions on this application server.

Any transaction that is not requested to complete before this time-out is rolled back. If set to 0, there is no time-out limit.

<b>Data type</b>	Integer
<b>Units</b>	Seconds
<b>Default</b>	120
<b>Range</b>	0 to 2 147 483 647

### Client inactivity timeout

Specifies the maximum duration, in seconds, between transactional requests from a remote client.

Any period of client inactivity that exceeds this timeout results in the transaction rolling back in this application server. If set to 0, there is no timeout limit.

<b>Data type</b>	Integer
<b>Units</b>	Seconds
<b>Default</b>	60
<b>Range</b>	0 to 2 147 483 647

### Enable logging for heuristic reporting

Select this property to enable the application server to log "about to commit one-phase resource" events from transactions that involve a one-phase commit resource and two-phase commit resources.

This property enables logging for heuristic reporting. If applications are configured to allow one-phase commit resources to participate in two-phase commit transactions, reporting of heuristic outcomes that occur at application server failure requires extra information to be written to the transaction log. If enabled, one additional log write is performed for any transaction that involves both one- and two-phase commit resources. No additional records are written for transactions that do not involve a one-phase commit resource.

<b>Data type</b>	String
------------------	--------

**Default  
Range**

Cleared

**Cleared**

The application server does not log "about to commit one-phase resource" events from transactions that involve a one-phase commit resource and two-phase commit resources.

**Selected**

The application server does log "about to commit one-phase resource" events from transactions that involve a one-phase commit resource and two-phase commit resources.

## Maximum Transaction Timeout

Specifies the maximum duration, in seconds, that transactions started by or propagated into this application server are allowed to execute.

**Data type**

Integer

**Units**

Seconds

**Default**

300

**Range**

0 to 2 147 040

---

## Using local transactions

Local transaction containment (LTC) support, and its configuration through local transaction extended deployment descriptors, gives IBM WebSphere Application Server application programmers a number of advantages. This topic describes those advantages and how they relate to the settings of the local transaction extended deployment descriptors. This topic also describes points to consider to help you best configure transaction support for some example scenarios that use local transactions.

### **Develop an enterprise bean or servlet that accesses one or more databases that are independent and require no coordination.**

If an enterprise bean does not need to use global transactions, it is often more efficient to deploy the bean with the Container Transaction deployment descriptor **Transaction** attribute set to Not supported instead of Required.

With the extended local transaction support of IBM WebSphere Application Server, applications can perform the same business logic in an unspecified transaction context as they can under a global transaction. An enterprise bean, for example, runs under an unspecified transaction context if it is deployed with a **Transaction** attribute of Not supported or Never.

The extended local transaction support provides a container-managed, implicit local transaction boundary within which application updates can be committed and their connections cleaned up by the container. Applications can then be designed with a greater degree of independence from deployment concerns. This makes using a **Transaction** attribute of Supports much simpler, for example, when the business logic may be called either with or without a global transaction context.

An application can follow a get-use-close pattern of connection usage regardless of whether or not the application runs under a transaction. The application can depend on the close behaving in the same way and not causing a rollback to occur on the connection if there is no global transaction.

There are many scenarios where ACID coordination of multiple resource managers is not needed. In such scenarios running business logic under a **Transaction** policy of Not supported performs better than if it had been run under a Required policy. This benefit is exploited through the **Local Transactions - Resolution-control** extended deployment setting of ContainerAtBoundary. With this setting, application interactions with resource providers (such as databases) are managed within implicit RMLTs that are both started and ended by the container. The RMLTs are committed

by the container at the configured **Local Transactions - Boundary**; for example at the end of a method. If the application returns control to the container by an exception, the container rolls back any RMLTs that it has started.

This usage applies to both servlets and enterprise beans.

#### **Use local transactions in a managed environment that guarantees clean-up.**

Applications that want to control RMLTs, by starting and ending them explicitly, can use the default **Local Transactions - Resolution-control** extended deployment setting of `Application`. In this case, the container ensures connection cleanup at the boundary of the local transaction context.

J2EE specifications that describe application use of local transactions do so in the manner provided by the default setting of **Local Transactions - Resolution-control**=`Application` and **Local Transactions - Unresolved-action**=`Rollback`. By configuring the **Local Transactions - Unresolved-action** extended deployment setting to `Commit`, then any RMLTs started by the application but not completed when the local transaction containment ends (for example, when the method ends) are committed by the container. This usage applies to both servlets and enterprise beans.

#### **Extend the duration of a local transaction beyond the duration of an EJB component method.**

The J2EE specifications restrict the use of RMLTs to single EJB methods. This restriction is because the specifications have no scoping device, beyond a container-imposed method boundary, to which an RMLT can be extended. In WBI Server Foundation, you can exploit the **Local Transactions - Boundary** extended deployment setting to give the following advantages:

- Significantly extend the use-cases of RMLTs
- Make conversational interactions with one-phase resource managers possible through `ActivitySession` support.

**5.1 +** An `ActivitySession` is a WBI Server Foundation programming model extension that provides a distributed context with a boundary that is longer than a single method. You can extend the use of RMLTs over the longer `ActivitySession` boundary, which can be controlled by a client. The `ActivitySession` boundary reduces the need to use distributed transactions where ACID operations on multiple resources are not needed. This benefit is exploited through the **Local Transactions - Boundary** extended deployment setting of `ActivitySession`. Such extended RMLTs can remain under the control of the application or be managed by the container depending on the use of the **Local Transactions - Resolution-control** deployment descriptor setting.

#### **Coordinate multiple one-phase resource managers.**

For resource managers that do not support XA transaction coordination, a client can exploit `ActivitySession`-bounded local transaction contexts. Such contexts give a client the same ability to control the completion direction of the resource updates by the resource managers as the client has for transactional resource managers. A client can start an `ActivitySession` and call its entity beans under that context. Those beans can perform their RMLTs within the scope of that `ActivitySession` and return without completing the RMLTs. The client can later complete the `ActivitySession` in a commit or rollback direction and cause the container to drive the `ActivitySession`-bounded RMLTs in that coordinated direction.

To determine how best to configure the transaction support for an application, depending on what you want to do with transactions, consider the following points.

#### **General points**

- You want to start and end global transactions explicitly in the application (BMT session beans and servlets only).

For a session bean, set the **Transaction type** to `Bean` (to use bean-managed transactions) in the component's deployment descriptor. (You do not need to do this for servlets.)

- You want to access only one XA or non-XA resource in a method.

In the component's deployment descriptor, set **Local Transactions - Resolution-control** to `ContainerAtBoundary`. In the Container transaction deployment descriptor, set **Transaction** to `Supports`.

- You want to access several XA resources atomically across one or more bean methods.



In the Container transaction deployment descriptor, set **Transaction** to Required, Requires new, or Mandatory.

- You want to access several non-XA resource in a method without having to worry about managing your own local transactions.

In the component's deployment descriptor, set **Local Transactions - Resolution-control** to ContainerAtBoundary. In the Container transaction deployment descriptor, set **Transaction** to Not supported.

- You want to access several non-XA resources in a method and want to manage them independently.

In the component's deployment descriptor, set **Local Transactions - Resolution-control** to Application and set **Local Transactions - Unresolved-action** to **Rollback**. In the Container transaction deployment descriptor, set **Transaction** to Not supported.

#### **Points specific to WBI Server Foundation**

- You want to access one of more non-XA resources across multiple EJB method calls without having to worry about managing your own local transactions.

In the component's deployment descriptor, set **Local Transactions - Resolution-control** to ContainerAtBoundary, **Local Transactions - Boundary** to ActivitySession, and **Bean Cache - Activate at** to ActivitySession. In the Container transaction deployment descriptor, set **Transaction** to Not supported and set **ActivitySession** attribute to Required, Requires new, or Mandatory.

- You want to access several non-XA resources across multiple EJB method calls and want to manage them independently.

In the component's deployment descriptor, set **Local Transactions - Resolution-control** to Application, **Local Transactions - Boundary** to ActivitySession, and **Bean Cache - Activate at** to ActivitySession. In the Container Transaction deployment descriptor, set **Transaction** to Not supported and set **ActivitySession** attribute to Required, Requires new, or Mandatory.

- You want to use a single non-XA resource and one or more XAResources.

Use the Last Participant Support of WBI Server Foundation.

---

## **Managing active transactions**

Use this task to manage transactions that are active on an application server.

You can use this task to display a snapshot of all the transactions currently running on an application server. For each transaction, the following properties are shown: its local ID, global ID, and current status. The transaction status is shown as an integer value. The values correspond to the following status:

- 0 - active
- 1 - marked for rollback
- 2 - prepared
- 3 - committed
- 4 - rolled back
- 5 - unknown
- 6 - none
- 7 - preparing
- 8 - committing
- 9 - rolling back

You can also choose to finish transactions manually.

Under normal circumstances, transactions should run and complete (commit or rollback) automatically, without the need for intervention. However, in some circumstances, you may need to finish a transaction manually. For example, you may want to finish a transaction that has become stuck polling a resource manager that you know will not become available again within the desired timeframe.



**Note:** If you choose to finish a transaction on an application server, it is recorded as having completed in the transaction service logs for that server, so will not be eligible for recovery during server start up. If you finish a transaction, you are responsible for cleaning up any in-doubt transactions on the resource managers affected.

To manage the active transactions for an application server, use the administrative console to complete the following steps:

1. In the navigation pane, select **Servers-> Manage Application Servers** This displays a list of application servers in the content pane.
2. In the content pane, click *your\_app\_server* This displays the properties of the application server, *your\_app\_server*.
3. In the content pane, click the **Runtime** tab. This displays the runtime properties of the application server.
4. In the Additional Properties table, select **Transaction Service** This displays the runtime properties of the Transaction Service.
5. Click **Manage Transactions**. This displays a snapshot of all the transactions currently running on the server. For each transaction, the following properties are shown: its local ID, current status, and global ID.
6. If you want to finish one or more transactions, select the checkbox provided on the entry for the transaction, then click **Finish**. Alternatively, to finish all transactions, select the checkbox in the header of the transactions table, then click **Finish**.

---

## Managing transaction logging for optimum server availability

This topic describes some considerations and actions that you can use to manage transaction logging to help ensure that the availability of your application servers is optimized.

The transaction service writes information to the transaction log for every global transaction which involves two or more resources or is distributed across multiple servers. The transaction log is stored on disk and is used by the transaction service for recovery after a system or server crash. The transaction log for each application server consists of multiple subdirectories and files held in a single directory. You can change the directory that an application server uses to store the transaction log, as described in Configuring transaction properties for an application server.

When a global transaction is completed, the information in the transaction log is not needed anymore so is marked for deletion. Periodically, this redundant information is garbage collected and the space reused by new transactions. The log files are created of fixed size at server startup, thus no further disk space allocation is required during the lifetime of the server. The default allocation is suitable for around 4000 concurrent transactions.

If all the log space is in use when a transaction needs to save information, the transaction is rolled back and the message WTRN0083W: The transaction log is full. Transaction rolled back. is reported to the system error log. No more transactions can commit until more log space is made available when existing active transactions complete.

You can monitor the number of concurrent global transactions by using the performance monitoring counters for transactions. The “Global transaction commit time” counter is a measure of how long a transaction takes to complete and, therefore, how long the log is in use by a transaction. If this value is high, then transactions are taking a long time to complete, which can be due to resource manager or network failures. If you ensure this value is low, the log is more efficiently used and unlikely to become full.

You can change the default size of log files by updating the transaction log settings as described in Configuring transaction properties for an application server.

## Configuring transaction aspects of servers for optimum availability

This topic describes some considerations and actions that you can take to configure transaction-related aspects of application servers for optimum availability.

To configure transaction-related aspects of application servers for optimum availability, complete the following steps:

1. Store the transaction log files on a fast disk in a highly-available file system, such as a RAID device. The transaction log may need to be accessed by every global transaction and be used for transaction recovery after a crash. Therefore, the disk the log files are being written to should be on a highly-available file system, such as a RAID device.

The performance of the disk also directly affects the transaction performance. In general, a global transaction makes two disk writes, one after the prepare phase when the outcome of the transaction is known (this information is forced to disk) and a further disk write at transaction completion. Therefore, the transaction logs should be placed on the fastest disks available and not make use of network mounted devices.

2. Mirror the transaction log files by using hardware disk mirroring or dual-ported disks. If log files have been mirrored or can be recovered, they can be used when restarting a failed server or moved to another machine and another server started there to perform recovery.

Hardware disk mirroring or dual-ported disks can be used by specifying the appropriate file system directory for the transaction logs using the WebSphere Administrative Console.

3. Specify the optimum location of the transaction log directory for application servers. By default, an application server places transaction log files in a subdirectory of the installed WebSphere Application Server, where the subdirectory name is the same as the server name. For example, the default directory for an application server named `server1` on Windows is `c:\WebSphere\AppServer\tranlog\server1`.

**Note:** The `TRANLOG_ROOT` variable was used in previous versions of WebSphere Application Server to override the default location of the transaction log, but is now deprecated. In a future version of WebSphere Application Server, transaction and compensation logs will move to a different default location and the variable `TRANLOG_ROOT` will be removed.

You can define a specific location for the transaction log directory for an application server by setting the **Transaction Log Directory** property for the server.

4. Never allow more than one application server to concurrently use the same set of log files. Because the transaction logs record the state of global transactions within a server, if the logs become lost or corrupt, then transactions that are in the prepared state before failure can leave resources in an in-doubt state and prevent further updates or access to the resources by other users or servers. These transactions may need to be manually resolved by either committing or rolling back the transactions at the affected resource managers. The failed server can then be cold-started, which creates new empty transaction logs.

If log files have been mirrored or can be recovered, they can be used when restarting the failed server or moved to an alternate server or machine and another server restarted to perform recovery, as described in the related tasks.

Never allow more than one application server to concurrently use the same set of log files, because each server will destroy the information recorded by the other, resulting in corrupt log files that are unusable for future recovery purposes.

5. Configure application servers to always use the same listening port address at each startup. If you are running distributed transactions between multiple application servers, the remote object references saved in the transaction log need to be redirected to the originating server on recovery.

On Application Server Network Deployment, the node agents automatically redirect such remote object references to the appropriate application servers on recovery. However, if the distributed transaction is between application servers that are not on Application Server Network Deployment, then you must handle the redirection of remote object references for transaction recovery to complete. For example,

you must do this is if an application server is deployed on WebSphere Application Server (not the Network Deployment edition) and runs distributed transactions with non-WebSphere EJB or Corba servers.

In particular, the default restart action of an application server not on Application Server Network Deployment is to use a different listening port address to the port when the server shut down. This prevents transaction recovery completing. To overcome this, you should always configure application servers to always use the same listening port address at each startup. See the ORB property `com.ibm.CORBA.ListenerPort` in ORB service settings that can be added to the administrative console. (See "Object Request Broker service settings" in the online help or Information Center.) You may need to make similar configuration changes to other application servers involved in transactions, to be able to access those servers during recovery.

## Moving a transaction log from one server to another

This topic describes some considerations and actions that you can take to move the transaction logs for an application server to another server.

To move transaction logs from one application server to another, consider the following steps:

1. Move all the transaction log files for the application server.

**5.1 +** The transaction log directory for each server contains a number of files and subdirectories. When moving transaction logs from one server to another you must move all of the files and subdirectories together as a single unit; otherwise recovery may not complete resulting in data inconsistency.

2. For a server configuration where there are no distributed transactions, move the transaction logs to any server that has access to the same resource managers. For a single server or network-deployed server configuration where it is known there are no distributed transactions present in the logs, the transaction logs can be moved to any server (on any node) that has access to the same resource managers as the original server. For example, the server needs communication and valid security access to databases or message queues.

All the transaction log files for the original server need to be moved to a directory accessible by the new server. This can be accomplished by either renaming the transaction log directory or copying all the contents to the new server's transaction log directory before starting the new server.

**Note:** To complete transaction recovery, the application server uses the resource manager configuration information in the transaction logs. However, for the application server to continue to do new work with the same resource managers, the server must have an appropriate resource manager configuration (as for the original server).

3. For a network-deployed server configuration where there are distributed transactions, move the transaction logs to a server that has the same name and host IP address, and access to the same resource managers. For a network-deployed server configuration, when it is known there are distributed transactions present in the logs, there are more restrictions. Distributed transactions that access multiple servers log information about each server involved in the transaction. This information includes the server name and the IP address of the machine on which the server is running. When recovery is taking place on server restart, the server uses this information to contact the distributed servers and similarly, the distributed servers try to contact the server with the same original name. So, if a server fails and the logs need to be recovered on an alternative server, that alternative server needs to have the same name and host IP address as the original server. The alternative server also needs to have access to the same resource managers as the original server. For example, the server needs communication and valid security access to databases or message queues.

**Note:** All servers within a cell must have unique names.

**Note:** To complete transaction recovery, the application server uses the resource manager configuration information in the transaction logs. However, for the application server to continue

to do new work with the same resource managers, the server must have an appropriate resource manager configuration (as for the original server).

## Restarting an application server on a different host

This topic describes some considerations and actions that you can take with transaction logs to restart an application server on a different host.

Moving transactions logs to a different host is similar to moving logs from one server to another, as described in *Moving transaction logs from one server to another*.

This involves moving an original application server on one host to an alternative server, which has access to the same resource managers, on another host. For a network-deployed server configuration, the alternative server must have the same name and host IP address as the original server.

**Note:** To complete transaction recovery, the application server uses the resource manager configuration information in the transaction logs. However, for the application server to continue to do new work with the same resource managers, the server must have an appropriate resource manager configuration (as for the original server).

To restart an application server on a different host, complete the following steps:

1. Ensure that the alternative application server is stopped.
2. Move all the transaction logs for the original server to the alternative application server, according to the considerations described in *Moving transaction logs from one server to another*.
3. Restart the alternative application server.

---

## Interoperating transactionally between application servers

This topic describes some considerations and actions that you can take to interoperate transactionally between different types of application servers.

To interoperate transactionally with a non-WebSphere application server, WebSphere Application Server switches dynamically between WebSphere-optimized native transaction contexts and interoperable OTS contexts depending on the capability of the partner with which it is interoperating.

**5.1+** To interoperate transactionally between WebSphere Application Server version 5.1 (or later) and WebSphere Application Server before version 5.0, you need to set the following system properties on application servers before version 5.0:

```
com.ibm.ejs.jts.jts.ControlSet.nativeOnly=false
com.ibm.ejs.jts.jts.ControlSet.interoperabilityOnly=true
```

For example, if you want to interoperate between application servers at WebSphere Application Server version 5.1 and WebSphere Application Server version 4.0.n, you need to set the system properties on the version 4.0.n application servers.

---

## Troubleshooting transactions

Use this overview task to help resolve a problem that you think is related to the Transaction service.

To identify and resolve transaction-related problems, you can use the standard WebSphere Application Server RAS facilities. If you encounter a problem that you think might be related to transactions, complete the following steps:

1. Check for transaction messages in the administrative console. The Transaction service produces diagnostic messages prefixed by “WTRN”. The error message indicates the nature of the problem and provides some detail. The associated message information provides an explanation and any user actions to resolve the problem.
2. Check for Transaction messages in the activity log. Activity log messages produced by the Transaction service are accompanied by Log Analyzer descriptions.
3. Check for more messages in the application server’s stdout.log. For more information about a problem, check the stdout.log file for the application server, which should contain more error messages and extra details about the problem.
4. Check for messages related to the application server’s transaction log directory when the problem occurred.

**Note:** If you changed the transaction log directory and a problem caused the application server to fail (with in-flight transactions) before the server was restarted properly, the server will next start with the new log directory and be unable to automatically resolve in-flight transactions that were recorded in the old log directory. To resolve this, you can copy the transaction logs to the new directory then stop and restart the application server.

---

## Transaction service exceptions

This topic lists the exceptions that can be thrown by the WebSphere Application Server transaction service. The exceptions are listed in the following groups:

- Standard exceptions
- Heuristic exceptions

If the EJB container catches a system exception from the business method of an enterprise bean, and the method is running within a container-managed transaction, the container rolls back the transaction before passing the exception on to the client. For more information about how the container handles the exceptions thrown by the business methods for beans with container-managed transaction demarcation, see the section *Exception handling* in the Enterprise JavaBeans 2.0 specification. That section specifies the container’s action as a function of the condition under which the business method executes and the exception thrown by the business method. It also illustrates the exception that the client receives and how the client can recover from the exception.

### Standard exceptions

The standard exceptions such as `TransactionRequiredException`, `TransactionRolledbackException`, and `InvalidTransactionException` are defined in the Java Transaction API (JTA) 1.0.1 Specification.

#### **InvalidTransactionException**

This exception indicates that the request carried an invalid transaction context.

#### **TransactionRequiredException exception**

This exception indicates that a request carried a null transaction context, but the target object requires an active transaction.

#### **TransactionRolledbackException exception**

This exception indicates that the transaction associated with processing of the request has been rolled back, or marked for roll back. Thus the requested operation either could not be performed or was not performed because further computation on behalf of the transaction would be fruitless.

### Heuristic exceptions

A heuristic decision is a unilateral decision made by one or more participants in a transaction to commit or rollback updates without first obtaining the consensus outcome determined by the Transaction Service. Heuristic decisions are an issue only after the participant has been prepared and the second phase of commit processing is underway. Heuristic decisions are normally made only in unusual circumstances, such as repeated failures by the transaction manager to communicate with a resource manager during

two-phase commit. If a heuristic decision is taken, there is a risk that the decision differs from the consensus outcome, resulting in a loss of data integrity.

The following list provides a summary of the heuristic exceptions. For more detail, see the Java Transaction API (JTA) 1.0.1 Specification.

**HeuristicRollback exception**

This exception is raised on the commit operation to report that a heuristic decision was made and that all relevant updates have been rolled back.

**HeuristicMixed exception**

This exception is raised on the commit operation to report that a heuristic decision was made and that some relevant updates have been committed and others have been rolled back.

---

## UserTransaction interface - methods available

For details about the methods available with the UserTransaction interface, see the WebSphere Application Server application programming interface reference information (Javadoc) or the Java Transaction API (JTA) 1.0.1 Specification.

---

## Using one-phase and two-phase commit resources in the same transaction

Use these topics to help you coordinate the use of a single one-phase commit capable resource with any number of two-phase commit capable resources in the same global transaction.

You can coordinate the use of a single one-phase commit capable resource with any number of two-phase commit capable resources in the same global transaction.

At transaction commit, the two-phase commit resources are prepared first using the two-phase commit protocol, and if this is successful the one-phase commit-resource is then called to `commit(one_phase)`. The two-phase commit resources are then committed or rolled back depending on the response of the one-phase commit resource.

For more information about using one-phase and two-phase commit resources within the same transaction, see the following topics:

- Coordinating use of one-phase and two-phase commit resources within the same transaction
- **5.1+** Assembling an application to use one-phase and two-phase commit resources in the same transaction (using WebSphere Studio or the Assembly Toolkit).
- Configuring an application server to allow logging for heuristic reporting

## Coordinating access to 1-PC and 2-PC-capable resources within the same transaction

Last participant support enables the use of a single one-phase commit capable resource with any number of two-phase commit capable resources in the same global transaction.

At transaction commit, the two-phase commit resources are prepared first using the two-phase commit protocol, and if this is successful the one-phase commit-resource is then called to `commit(one_phase)`. The two-phase commit resources are then committed or rolled back depending on the response of the one-phase commit resource.

**Note:** **5.1+** If the global transaction is distributed across multiple application servers *that are all running at WebSphere Application Server version 5.1 or later*, you can coordinate access to one-phase and two-phase commit capable resources within the same transaction.



**Note:** If the global transaction is distributed across multiple application servers *that are all running at WebSphere Application Server version 5.1 or later* then you can exploit last participant support to coordinate a one-phase commit capable resource and any number of two-phase commit capable resources within the same transaction, in a limited number of scenarios.

- The main scenario is where the one-phase commit resource provider is accessed in the application server process (the “transaction root” server) in which the transaction is started.  
In this scenario, last participant support can coordinate a one-phase commit capable resource and any number of two-phase commit capable resources within the same transaction.
- If the one-phase commit resource provider is accessed in a different application server (a “transaction subordinate” server) from the one in which the transaction was started; for example, as a result of a transactional invocation on a remote EJB interface where the EJB implementation accesses a one-phase commit resource provider.

In this scenario, the transaction typically cannot be committed. To be able to commit (as part of a global transaction) a one-phase commit resource enlisted on a transaction subordinate server, the transaction service must delegate coordination responsibility from the transaction root to the subordinate server. This occurs only if no other resources were registered with the transaction root server.

Last participant support introduces an increased risk of an heuristic outcome to the transaction. That is, the transaction manager cannot be sure that all resources were completed in the same direction (either committed or rolled back). For this reason, to enable an application to coordinate access to one-phase and two-phase commit capable resources within the same transaction, you configure the application to accept the increased risk of an heuristic outcome.

An heuristic outcome occurs if the transaction service (JTS) receives no response from the commit one-phase flow on the one-phase commit resource. In this situation the transaction service cannot determine whether changes for the one-phase commit resource were committed or rolled back, so cannot drive reliably the correct outcome of the global transaction on the other two-phase commit resources.

You can configure the transaction service for an application server to indicate whether or not to log that it is about to commit the one-phase commit resource. This does not reduce the heuristic hazard, but ensures that any failure, and subsequent recovery, of the application server during the one-phase commit phase occurs with knowledge of whether or not the one-phase commit resource was asked to commit:

- If the one-phase commit resource was asked to commit, a heuristic outcome is reported to the activity log.
- If the one-phase commit resource was not asked to commit, then the transaction is rolled back consistently.

## **Assembling an application to use one-phase and two-phase commit resources in the same transaction**

Use this task to assemble an application to use one-phase and two-phase commit resources within the same transaction.

To enable an application to use one-phase and two-phase commit capable resources within the same transaction, you must configure the deployment attributes of the application to accept the increased risk of an heuristic outcome.

You can configure the deployment attributes of an application by using the Deployment Descriptor Editor of WebSphere Studio Application Developer or the Assembly Toolkit.

This task description assumes that you have an EAR file for an application component, that can be deployed in WebSphere Application Server. For more details about using the Assembly Toolkit, see *Assembling applications with the Assembly Toolkit*.



To configure an application to indicate that you accept the increased risk of an heuristic outcome, complete the following steps:

1. Start the Assembly Toolkit.
2. Create or edit the application EAR file.

**Note:** Ensure that you set the target server as WBI Server Foundation v5.1.

For example, to change attributes of an existing application, use the import wizard to import the EAR file into the Assembly Toolkit. To start the import wizard:

- a. Click **File-> Import-> EAR file**
  - b. Click **Next**, then select the EAR file.
  - c. In the Target server field, select WBI Server Foundation v5.1
  - d. Click **Finish**
3. In the J2EE Hierarchy view, right-click the Enterprise Application instance, then click **Open With > Deployment Descriptor Editor**. A property dialog notebook for the component is displayed in the property pane.
  4. In the property pane, select the Extended Services tab.
  5. In the Last Participant Support section, select the **Last participant support** checkbox.
  6. Save your changes to the deployment descriptor.
    - a. Close the Deployment Descriptor Editor.
    - b. When prompted, click **Yes** to indicate that you want to save changes to the deployment descriptor.
  7. Verify the archive files.
  8. Generate code for deployment for EJB modules or for enterprise applications that use EJB modules.
  9. Optional: Test your completed module on a WebSphere Application Server installation. Right-click a module, click **Run on Server**, and follow the instructions in the displayed wizard. Note that **Run on Server** works on the Windows, Linux/Intel, and AIX operating systems only; you cannot deploy remotely from the Assembly Toolkit to a WebSphere Application Server installation on a UNIX operating system such as Solaris.

### **Important**

**Important:** Use **Run On Server** for unit testing only. Assembly Server Toolkit controls the WebSphere Application Server installation and, when an application is published remotely, the Toolkit overwrites the server configuration file for that server. Do not use on production servers.

For instructions on remote testing, see the article "Setting Up a Remote WebSphere Application Server in WebSphere Studio V5" at [http://www7b.boulder.ibm.com/wsdd/techjournal/0303\\_yuen/yuen.html](http://www7b.boulder.ibm.com/wsdd/techjournal/0303_yuen/yuen.html).

After assembling your application, use a systems management tool to deploy the EAR file onto the application server that is to run the application; for example, using the administrative console as described in Deploying and managing applications.

## **Last participant support extension settings**

Use this page to configure last participant support extensions.

Last participant support is an extension to the transaction service to allow a single one-phase resource to participate in a two-phase transaction with one or more two-phase resources.

To view this administrative console page, click **Applications > Applications > *application\_name* > Last Participant Support Extension**.

### **Accept Heuristic Hazard:**

Specifies whether the application accepts the possibility of an heuristic hazard occurring in a two-phase transaction containing a one-phase resource.

**Default  
Range**

Cleared  
**Selected**

The application accepts the increased risk of an heuristic outcome.

**Cleared**

The application does not accept the increased risk of an heuristic outcome.

## Configuring an application server to log heuristic reporting

To enable an application server to log “about to commit one-phase resource” events from transactions that involve a one-phase commit resource and two-phase commit resources, use the Administrative console to complete the following steps:

1. Start the Administrative console
2. In the navigation pane, select **Servers-> Manage Application Servers-> *your\_app\_server*** This displays the properties of the application server, *your\_app\_server*, in the content pane.
3. Select the Transaction Service tab, to display the properties page for the transaction service, as two notebook pages:

### **Configuration**

The values of properties defined in the configuration file. If you change these properties, the new values are applied when the application server next starts.

### **Runtime**

The runtime values of properties. If you change these properties, the new values are applied immediately, but are overwritten with the Configuration values when the application server next starts.

4. Select the Configuration tab, to display the transaction-related configuration properties.
5. Select the **Enable logging for heuristic reporting** checkbox.
6. Click **OK**.
7. Stop then restart the application server.

## Exceptions thrown for transactions involving both single- and two-phase commit resources

The exceptions that can be thrown by transactions that involve single- and two-phase commit resources are the same as those that can be thrown by transactions involving only two-phase commit resources.

The exceptions that can be thrown are listed in the WebSphere API reference information (Javadoc).

## Last Participant Support: Resources for learning

Use the links in this topic to find relevant supplemental information about Last Participant Support. The information resides on IBM and non-IBM Internet sites, whose sponsors control the technical accuracy of the information.

These links are provided for convenience. Often, the information is not specific to the IBM WebSphere Application Server product, but is useful all or in part for understanding the product. When possible, links are provided to technical papers and Redbooks that supplement the broad coverage of the release documentation with in-depth examinations of particular product areas.

View links to additional information about:

- “Programming specifications” on page 1275
- “Other” on page 1275

**Programming specifications**

- J2EE Activity Service for Extended Transactions
- Java Transaction API (JTA) 1.0.1

**Other**

- WebSphere Application Server Enterprise Version 5 Overview: Advanced Transactional Connectivity
- Listing of PDF files to learn about WebSphere Application Server Version 5
- Listing of all IBM WebSphere Application Server Redbooks
- Listing of all IBM WebSphere Application Server Whitepapers
- WebSphere Application Server Enterprise Edition 4.0: A Programmer's Guide



---

## Chapter 25. Using the ActivitySession service

These topics provide information about implementing WebSphere enterprise applications that use ActivitySessions.

The ActivitySession service provides an alternative unit-of-work (UOW) scope to that provided by global transaction contexts. ActivitySessions provide a scoping mechanism for units of work, and both an ActivitySession and a transaction has the same following characteristics:

- It can be bean-managed or container-managed
- It can be distributed across application servers
- It can be used as the context for managing EJB activation policy and lifecycle

An ActivitySession differs significantly from a transaction in the manner of its interaction with resource managers. An ActivitySession is used to scope or coordinate local transactions. That is, an ActivitySession can be used to request multiple one-phase resource managers to come to an application- or container-determined outcome. Unlike a transaction, an ActivitySession has no notion of a prepare phase or any notion of recovery at a service level.

The WebSphere EJB container and deployment tools support ActivitySessions as an extension to the J2EE programming model. Enterprise beans can be deployed with lifecycles that are influenced by ActivitySession context, as an alternative to transaction context. An enterprise bean with an ActivitySession-scoped lifecycle can participate in a resource manager local transaction (RMLT) that has a duration of the ActivitySession rather than an individual method on the bean (which is all that is possible under the standard J2EE model). Applications can then be composed of several enterprise beans with ActivitySession-based activation, with each bean participating in extended local transactions with one or more resource managers. At the end of the ActivitySession each of the local transactions can be directed to a common outcome by the ActivitySession manager.

You can configure the WebSphere containers and deployable applications to support enterprise beans that operate under application- or container-initiated ActivitySessions rather than, or in addition to, transactions.

For more information about implementing WebSphere enterprise applications that use ActivitySessions, see the following topics:

- The ActivitySession service
  - ActivitySessions and transaction contexts
  - Using ActivitySessions with HTTP sessions
- The ActivitySession service programming interfaces
- Developing a J2EE application to use an ActivitySession
- ActivitySessions samples
- Configuring EJB module ActivitySession deployment attributes with the Assembly Toolkit
- Configuring Web module ActivitySession deployment attributes with the Assembly Toolkit
- Disabling or enabling the ActivitySession service
- Configuring the default ActivitySession timeout
- Troubleshooting ActivitySessions

---

### The ActivitySession service

The ActivitySession service provides an alternative unit-of-work (UOW) scope to that provided by global transaction contexts. An ActivitySession context can be longer-lived than a global transaction context and can encapsulate global transactions.

Support for the ActivitySession service is shown in the following figure:

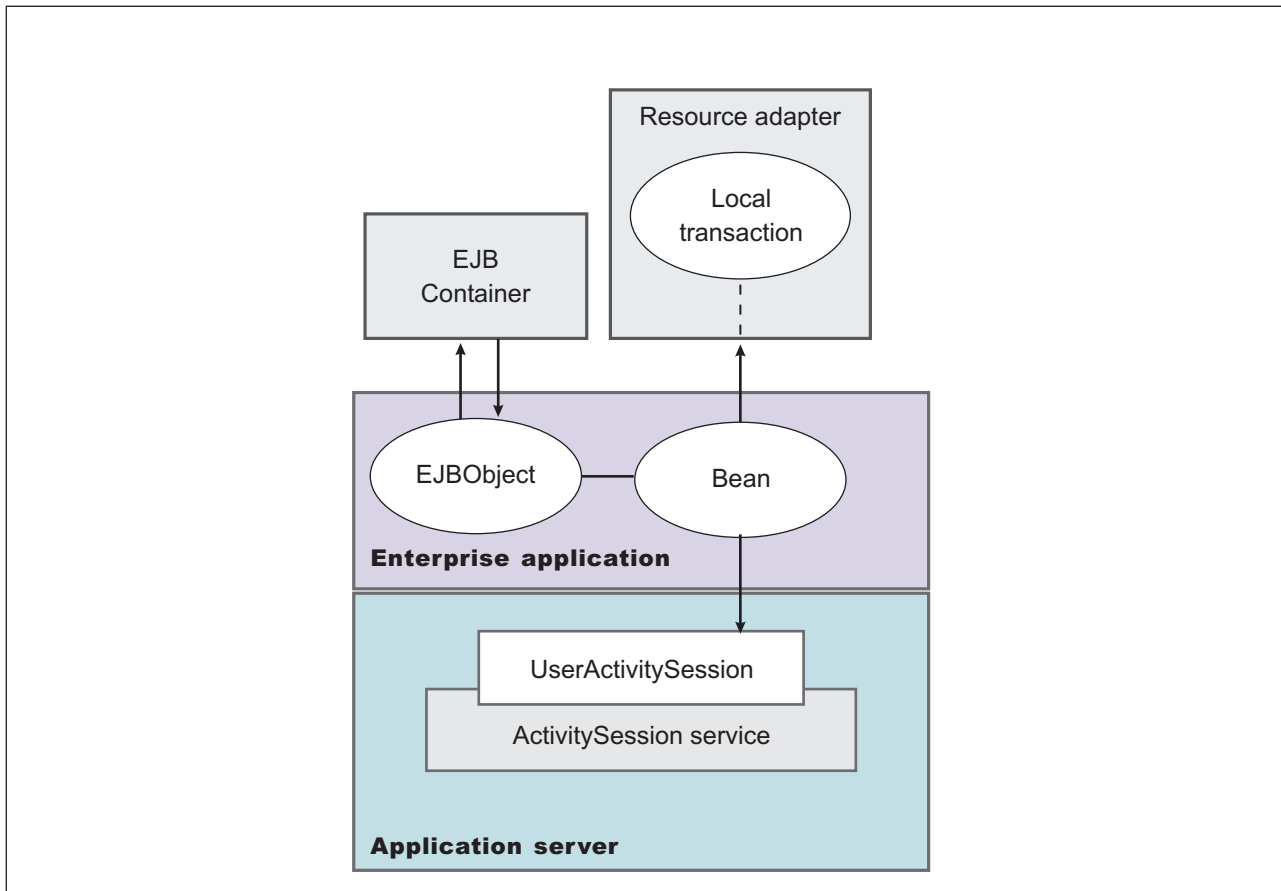


Figure 33. The ActivitySession service. This figure show the main components of the ActivitySession service within WebSphere Application server. For an overview of these components, see the text that accompanies this figure.

Although the purpose of a global transaction is to coordinate multiple resource managers, global transaction context is often used by J2EE applications as a “session” context through which to access EJB instances. An ActivitySession context is such a session context, and can be used in preference to a global transaction in cases where coordination of two-phase commit resource managers is not needed. Further, an ActivitySession can be associated with an HttpSession to extend a “client session” to an HTTP client.

ActivitySession support is available to Web, EJB, and J2EE-client components. EJB components can be divided into beans that exploit container-managed ActivitySessions and beans that use bean-managed ActivitySessions.

The ActivitySession service provides a UserActivitySession application programming interface available to J2EE components that use bean-managed ActivitySessions for application-managed demarcation of ActivitySession context. The ActivitySession service also provides a system programming interface for container-managed demarcation of ActivitySession context and for container-managed enlistment of one-phase resources (RMLTs) in such contexts.

The UserActivitySession interface is obtained by a JNDI lookup of `java:comp/websphere/UserActivitySession`. This interface is not available to enterprise beans that use container-managed ActivitySessions, and any attempt by such beans to obtain the interface results in a `NotFound` exceptions.

## Using ActivitySessions with HTTP sessions

This topic describes how a web application that runs in the WebSphere Web container can participate in an ActivitySession context.

If the web application is designed such that several servlet invocations occur as part of the same logical application, then the servlets can use the `HttpSession` to preserve state across servlet invocations. The `ActivitySession` context is one state that can be suspended into the `HttpSession` and resumed on a future invocation of a servlet that accesses the `HttpSession`.

An `ActivitySession` is associated automatically with an `HttpSession`, so can be used to extend access to the `ActivitySession` over multiple HTTP invocations, over inclusion or forwarding of servlets, and to support EJB activation periods that can be determined by the lifecycle of the web HTTP client. An `ActivitySession` context stored in an `HttpSession` can also be used to relate work for the `ActivitySession` back to a specific web HTTP client.

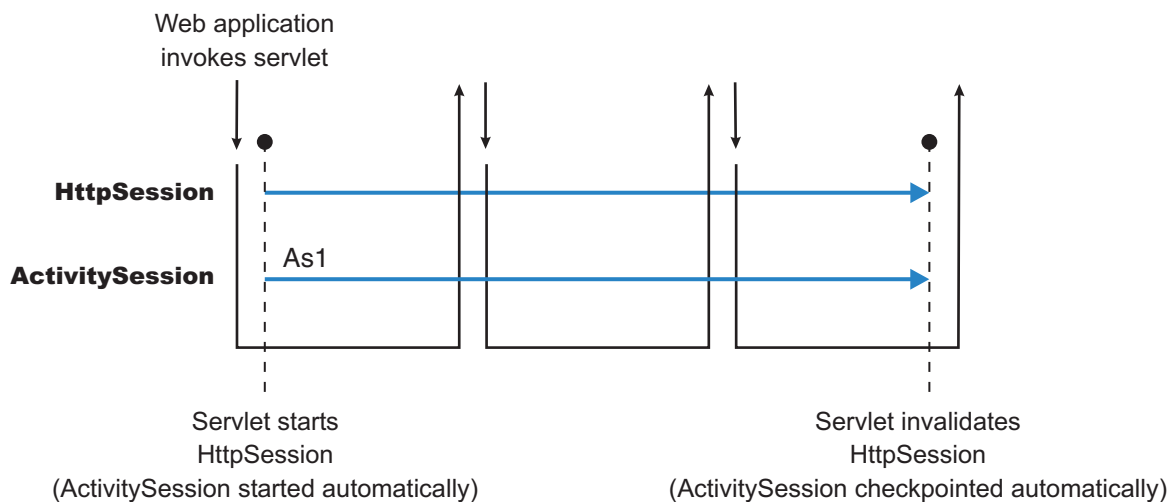
The Web container manages `ActivitySessions` based on deployment descriptor attributes associated with servlets in the Web application module. The two usage models are:

- The Web container starts and ends `ActivitySessions`.

The Web application invokes a servlet that has been configured for container control of `ActivitySessions`.

- If an `HttpSession` exists then it has an associated `ActivitySession`.
- If an `HttpSession` does not exist, the servlet can start an `HttpSession`, which causes an `ActivitySession` to be started automatically and associated with the `HttpSession`.

A servlet cannot start a new `HttpSession` until an existing `HttpSession` has been ended. Within an `HttpSession`, the Web application can invoke other servlets that can use the associated `ActivitySession` context. When the Web application invokes a servlet that ends the `HttpSession`, the `ActivitySession` is ended automatically. This is shown in the following diagram:



- The Web application starts and ends `ActivitySessions`.

The Web application invokes a servlet that has been configured for application control of `ActivitySessions`.

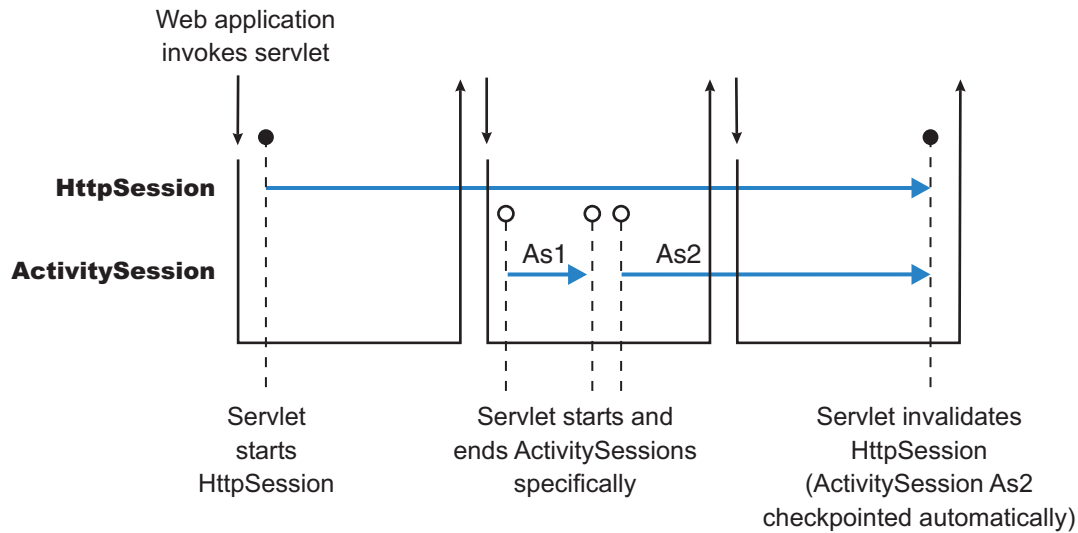
- If an `HttpSession` exists and has an associated `ActivitySession`, the servlet can use or end that `ActivitySession` context.
- If an `HttpSession` does not exist, the servlet can start an `HttpSession`, but this does not automatically start an `ActivitySession`.
- If an `HttpSession` exists but does not have an associated `ActivitySession`, the servlet can start a new `ActivitySession`. This automatically associates the `ActivitySession` with the `HttpSession`. The `ActivitySession` lasts either until the `ActivitySession` is specifically ended or until the `HttpSession` is ended.

The servlet cannot start a new `ActivitySession` until an existing `ActivitySession` has been ended. The servlet cannot start a new `HttpSession` until an existing `HttpSession` has been ended.

Within an `HttpSession`, the Web application can invoke other servlets that can use or end an existing `ActivitySession` context or, if no `ActivitySession` exists start a new `ActivitySession`. When the Web



application invokes a servlet that ends the HttpSession, the HttpSession is ended automatically. This is shown in the following diagram:



A Web application can invoke servlets configured for either usage model.

The following points apply to both usage models:

- To end an HttpSession (and any associated ActivitySession), the Web application must invalidate that session. This causes the ActivitySession to be checkpointed.
- Any downstream EJBs activated within the context of an ActivitySession can be held in memory rather than passivated between servlet invocations, because the client effectively becomes the web HTTP client.
- Web applications can be composed of many servlets, and each servlet in the Web application can be configured with a value for ActivitySessionControl. ActivitySessionControl determines whether the servlet or its container starts any ActivitySessions.
- An ActivitySession context that encapsulates an active transaction context cannot be associated with an HttpSession, because a transaction can hold database locks and should be designed to be shortlived. If an application moves an active transaction to an HttpSession, the transaction is rolled back and the ActivitySession is suspended into the HttpSession. In general, you should design applications to use ActivitySessions or other constructs as the long-lived entities and ACID transactions as short-duration entities within these.
- Only one ActivitySession can be associated with an HttpSession at any time, for the duration of the ActivitySession. An ActivitySession associated with an HttpSession remains associated for the duration of that ActivitySession, and cannot be replaced with another until the first ActivitySession is completed. The ActivitySession can be accessed by multiple servlets if they have shared access to the HttpSession.
- ActivitySessions are not persistent. If a persistent HttpSession exists longer than the server hosting it, any cached ActivitySession is terminated when the hosting server ends.
- If the HttpSession times out before the associated ActivitySession has ended, the ActivitySession is reset.
- If the ActivitySession times out, it is reset then the HttpSession is ended.

## ActivitySession and transaction contexts

This topic describes the hierarchical relationship between transaction and ActivitySession context, This relationship, defined by the ActivitySession service, requires that any transaction context be either wholly inside or wholly outside an ActivitySession context.

An `ActivitySession` context is very similar to a transaction context and extends the lifecycle choices for activation of enterprise beans; it can encapsulate one or more transactions. The `ActivitySession` context is a distributed context that, like the transaction context, can be bean- or container-managed. An `ActivitySession` context is used mainly by a client to scope the lifecycle of an enterprise bean that it uses either beyond or in the absence of individual transactions started by that client.

`ActivitySessions` have a lower overhead than transactions and can be used instead of transactions that are only used to scope the lifecycle of a called enterprise bean. For a bean with an activation policy of `ActivitySession`, the duration of any resource manager local transactions (RMLTs) started by that bean can be bounded by the duration of the `ActivitySession` instead of the bean method in which the RMLT was started. This provides flexibility and potential for using RMLTs in an enterprise bean beyond the scenarios described in the J2EE specifications. The J2EE specifications define that RMLTs need to be completed before the end of the bean method, because the bean method is the only containment boundary for local transactions available in those specifications.

The following rules defines the relationship between transactions and `ActivitySessions`.

- The EJB or Web container always uses a local transaction containment (LTC) if there is no global transaction present. An LTC can be method-scoped or `ActivitySession`-scoped.
- Before a method dispatch, the container ensures that there is always either an LTC or global transaction context, but never both contexts.
- `ActivitySessions` cannot be nested within each other. Any attempt to start a nested `ActivitySession` results in a `com.ibm.websphere.ActivitySession.NotSupportedException` on `UserActivitySession.beginSession()`.
- An `ActivitySession` can wholly encapsulate one or more global transactions.
- An `ActivitySession` cannot be encapsulated by a global transaction nor should `ActivitySession` and global transaction boundaries overlap. Any attempt to start an `ActivitySession` in the presence of a global transaction context results in a `com.ibm.websphere.ActivitySession.NotSupportedException` on `UserActivitySession.beginSession()`. Any attempt to call `endSession(EndModeCheckpoint)` on an `ActivitySession` that contains an incomplete global transaction results in a `com.ibm.websphere.ActivitySession.ContextPendingException`. Neither the global transaction nor the `ActivitySession` context are affected. If `endSession(EndModeReset)` is called then the `ActivitySession` is reset and the global transactions marked `rollback_only`.
- Each global transaction wholly encapsulated by an `ActivitySession` is independent of every other global transaction within that `ActivitySession`. A rollback of one global transaction does not affect any others or the `ActivitySession` itself.
- `ActivitySession` and global transaction contexts can coexist with an `ActivitySession` encapsulating one or more serially-executing global transactions.

## Combining transaction and `ActivitySession` container policies

This topic provides details about the relationship between the deployment descriptor properties that determine how the container manages `ActivitySession` boundaries.

If an enterprise bean uses `ActivitySessions`, how the EJB container manages `ActivitySession` boundaries when delegating a method invocation depends on both the **ActivitySession kind** and **Transaction attribute** deployment descriptor attributes configured for the enterprise bean. The following table lists the relationship between these two properties.

In each row, the final column describes the behavior that the EJB container takes with respect to global transaction and `ActivitySession` context, based on the following abbreviations:

**S<sub>n</sub>** An `ActivitySession`, where *n* indicates the `ActivitySession` instance.

**T<sub>n</sub>** A transaction, where *n* indicates the transaction instance.

In every case where the container does not start or leave a global transaction context associated with the thread, it starts (or obtains from the bean instance) a local transaction containment and associates that with the thread. The duration of the local transaction containment is determined by a combination of the

local-transaction boundary descriptor (configured as part of the application deployment descriptor, and not shown in the following table) and the presence or not of an ActivitySession context, as described in ActivitySessions and transaction contexts.

The rows highlighted in bold are not allowed.

Table 11. Container behavior for activitysession and transaction policies deployment settings

<b>Bean ActivitySession policy(ActivitySession kind)</b>	<b>Bean transaction policy(Transaction attribute)</b>	<b>Received contexts</b>	<b>Container behaviour</b>
Required	Required	None	Start S1, Start T1
S1	Start T1		
T1	Suspend T1, Start S1, Start T2		
S1, T1	No Action		
Requires new	None	Start S1, Start T1	
S1	Start T1		
T1	Suspend T1, Start S1, Start T2		
S1, T1	Suspend T1, Start T2		
Supports	None	Start S1	
S1	No Action		
T1	Suspend T1, Start S1		
S1, T1	No Action		
Not supported	None	Start S1	
S1	No Action		
T1	Suspend T1, Start S1		
S1, T1	Suspend T1		
Mandatory	None	Exception	
S1	Exception		
T1	Exception		
S1, T1	No action		
Never	None	Start S1	
S1	No Action		
T1	Suspend T1, Start S1		
S1, T1	Exception		
Requires new	Required	None	Start S1 + T1
S1	Suspend S1, Start S2 + T1		
T1	Suspend T1, Start S1 + T2		
S1 + T1	Suspend S1 + T1, Start S2 + T2		
Requires new	None	Start S1 + T1	
S1	Suspend S1, Start S2 + T1		
T1	Suspend T1, Start S1 + T2		

Table 11. Container behavior for activitysession and transaction policies deployment settings (continued)

Bean ActivitySession policy(ActivitySession kind)	Bean transaction policy(Transaction attribute)	Received contexts	Container behaviour
S1 + T1	Suspend S1 + T1, Start S2 + T2		
Supports	None	Start S1	
S1	Suspend S1, Start S2		
T1	Suspend T1, Start S1		
S1, T1	Suspend S1 + T1, Start S2		
Not supported	None	Start S1	
S1	Suspend S1, Start S2		
T1	Suspend T1, Start S1		
S1, T1	Suspend S1 + T1, Start S2		
<b>Mandatory</b>	<b>None</b>	<b>Exception</b>	
<b>S1</b>	<b>Exception</b>		
<b>T1</b>	<b>Exception</b>		
<b>S1, T1</b>	<b>Exception</b>		
Never	None	Start S1	
S1	Suspend S1, Start S2		
T1	Suspend T1, Start S1		
S1, T1	Suspend S1 + T1, Start S2		
Supports	Required	None	Start T1
S1	Start T1		
T1	No Action		
S1, T1	No Action		
Requires new	None	Start T1	
S1	Start T1		
T1	Suspend T1, Start T2		
S1, T1	Suspend T1, Start T2		
Supports	None	No Action	
S1	No Action		
T1	No Action		
S1, T1	No Action		
Not supported	None	No Action	
S1	No Action		
T1	Suspend T1		
S1, T1	Suspend T1		
<b>Mandatory</b>	<b>None</b>	<b>Exception</b>	
<b>S1</b>	<b>Exception</b>		
<b>T1</b>	<b>No Action</b>		

Table 11. Container behavior for activitysession and transaction policies deployment settings (continued)

Bean ActivitySession policy(ActivitySession kind)	Bean transaction policy(Transaction attribute)	Received contexts	Container behaviour
S1, T1	No Action		
Never	None	No Action	
S1	No Action		
T1	Exception		
S1, T1	Exception		
Not supported	Required	None	Start T1
S1	Suspend S1, Start T1		
T1	No Action		
S1, T1	Suspend S1 + T1, Start T2		
Requires new	None	Start T1	
S1	Suspend S1, Start T1		
T1	Suspend T1, Start T2		
S1, T1	Suspend S1 + T1, Start T2		
Supports	None	No Action	
S1	Suspend S1		
T1	No Action		
S1, T1	Suspend S1 + T1		
Not supported	None	No Action	
S1	Suspend S1		
T1	Suspend T1		
S1, T1	Suspend S1 + T1		
Mandatory	None	Exception	
S1	Exception		
T1	No Action		
S1,T1	Exception		
Never	None	No Action	
S1	Suspend S1		
T1	Exception		
S1, T1	Suspend S1 + T1		
Mandatory	Required	None	Exception
S1	Start T1		
T1	Exception		
S1, T1	No Action		
Requires new	None	Exception	
S1	Start T1		
T1	Exception		

Table 11. Container behavior for activitysession and transaction policies deployment settings (continued)

Bean ActivitySession policy(ActivitySession kind)	Bean transaction policy(Transaction attribute)	Received contexts	Container behaviour
S1, T1	Suspend T1, Start T2		
Supports	None	Exception	
S1	No Action		
T1	Exception		
S1, T1	No Action		
Not supported	None	Exception	
S1	No Action		
T1	Exception		
S1, T1	Suspend T1		
Mandatory	None	Exception	
S1	Exception		
T1	Exception		
S1, T1	No Action		
Never	None	Exception	
S1	No Action		
T1	Exception		
S1,T1	Exception		
Never	Required	None	Start T1
S1	Exception		
T1	No Action		
S1, T1	Exception		
Requires new	None	Start T1	
S1	Exception		
T1	Suspend T1, Start T2		
S1,T1	Exception		
Supports	None	No Action	
S1	Exception		
T1	No Action		
S1,T1	Exception		
Not supported	None	No Action	
S1	Exception		
T1	Suspend T1		
S1,T1	Exception		
Mandatory	None	Exception	
S1	Exception		
T1	No Action		

Table 11. Container behavior for activitysession and transaction policies deployment settings (continued)

Bean ActivitySession policy(ActivitySession kind)	Bean transaction policy(Transaction attribute)	Received contexts	Container behaviour
S1,T1	Exception		
Never	None	No Action	
S1	Exception		
T1	Exception		
S1,T1	Exception		
Bean managed	Bean managed	None	No Action
S1	Suspend S1		
T1	Suspend T1		
S1, T1	Suspend S1 + T1		

## Developing a J2EE application to use ActivitySessions

This topic provides an overview of the scenarios for which you would develop a J2EE application to use an ActivitySession.

The following common J2EE application scenarios make use of an ActivitySession:

- Developing a J2EE application to use one or more enterprise beans that are persisted to non-transactional datastores.

This scenario can be used by an application that needs to coordinate multiple one-phase resource managers; for example, for two or more entity EJBs whose persistence is delegated to LocalTransaction resource adapters.

In this scenario, the enterprise beans used by the application have an Activation policy of ActivitySession and a local transaction containment policy with a boundary of ActivitySession and resolution-control of ContainerAtBoundary. The synchronization of the EJB state data is synchronized, by the container, with the one-phase resource managers at ActivitySession completion and no application code is required to be aware of ActivitySession support.

- Developing a J2EE application in which an enterprise bean accesses a resource manager multiple times in different business methods.

This scenario can be used by an application that needs to extend a resource manager local transaction (RMLT) over several business methods of an enterprise bean instance.

In this scenario, the enterprise beans used by the application have an Activation policy of ActivitySession and a local transaction containment policy with a boundary of ActivitySession and resolution-control of Application. The application logic starts and ends the RMLTs, for example using the `javax.resource.cci.LocalTransaction` interface offered by a LocalTransaction Connector, but is not constrained to start and commit the LocalTransaction in the same method.

- Developing a J2EE client application to use an ActivitySession to scope EJB activation and load-balancing.

This scenario can be used by an application client that needs to access an entity bean instance several times in the same client session, either without needing to run under a transaction context, or with the need to run under a number of distinct and serially-executed transactions.

In this scenario, the enterprise beans used by the application client have an Activation policy of ActivitySession and a local transaction containment policy appropriate to the function of the enterprise bean. The J2EE client application can represent a period of user activity, for example a signon period, during which a number of interactions occur with one or more enterprise beans. If the J2EE client application begins an ActivitySession and invokes the enterprise beans within the scope of the UOW



represented by the `ActivitySession`, then the enterprise bean instances are activated by the container on the `ActivitySession` boundary and remain in the active state until passivated by the container at the end of the `ActivitySession`. Workload affinity management based on the `ActivitySession` is a platform quality of service. Global transactions can begin and end within the `ActivitySession`, if they are wholly encapsulated by the `ActivitySession` and run serially. EJB instances activated at the `ActivitySession` boundary remain active across the serial global transactions.

- Developing a Web application client to participate in an `ActivitySession` context.

A web application that runs in the WebSphere Web container can participate in an `ActivitySession` context. Web applications can use the `UserActivitySession` interface to begin and end an `ActivitySession` context. Also, the `ActivitySession` can be associated with an `HttpSession`, thereby extending access to the `ActivitySession` over multiple HTTP invocations and supporting EJB activation periods that can be determined by the lifecycle of the web HTTP client.

The Web container manages `ActivitySessions` based on deployment descriptor attributes associated with the Web application module.

#### **General considerations:**

- An application that is accessed under an `ActivitySession` context can receive a `javax.transaction.InvalidTransactionException RemoteException`, thrown by the EJB container when servicing any application method. This exception occurs when an instance of an enterprise bean that has an `ActivitySession`-based activation policy becomes involved with concurrent global and local transactions.
- To enable an enterprise bean to participate in an `ActivitySession` context and support `ActivitySession`-based operations, it must be configured with an `ActivationPolicy` of `ACTIVITY_SESSION`. A bean configured with `ActivationPolicy` of either `TRANSACTION` or `ONCE` cannot participate in `ActivitySession` context.
- A session bean can either use container-managed `ActivitySessions` or implement bean-managed `ActivitySessions`; entity beans can only use container-managed `ActivitySessions`. A bean is deployed to be bean-managed or container-managed with respect to `ActivitySession` management by setting its transaction type deployment attribute to be bean-managed or container-managed when deploying the enterprise bean. A bean that uses bean-managed transactions can use bean-managed `ActivitySessions`; a bean that uses container-managed transactions can use container-managed `ActivitySessions`.
- If you want a session bean or J2EE client to manage its own `ActivitySessions`, you must write the code that explicitly demarcates the boundaries of an `ActivitySession`, as described in *Developing an enterprise bean or J2EE client to manage ActivitySessions*.

For examples of using `ActivitySessions` in J2EE applications, see *ActivitySessions samples*.

---

## **Developing an enterprise bean or J2EE client to manage ActivitySessions**

Use this task to write the code needed by a session EJB or J2EE client application to manage an `ActivitySession`, based on the example code extract provided.

In most situations, an enterprise bean can depend on the EJB container to manage `ActivitySessions` within the bean. In these situations, all you need to do is set the appropriate `ActivitySession` attributes in the EJB module deployment descriptor, as described in *Configuring EJB module ActivitySession deployment attributes with the Assembly Toolkit*. Further, in general, it is practical to design your enterprise beans so that all `ActivitySession` management is handled at the enterprise bean level.

However, in some cases you may need to have a session bean or J2EE client participate directly in `ActivitySessions`. You then need to write the code needed by the session bean or J2EE client application to manage its own `ActivitySessions`.

**Note:** Session beans that use BMT and have an **Activate at** setting of Activity session can manage ActivitySessions. Entity beans cannot manage ActivitySessions; the EJB container always manages ActivitySessions within entity beans.

When preparing to write code needed by a session bean or J2EE client application to manage ActivitySessions, consider the points described in ActivitySessions and transaction contexts.

To write the code needed by a session EJB or J2EE client application to manage an ActivitySession, complete the following steps based on the example code extract below:

1. Get an initial context for the ActivitySession.
2. Get an implementation of the UserActivitySession interface, by a JNDI lookup of the URL `java:comp/websphere/UserActivitySession`. The UserActivitySession interface is used to begin and end ActivitySessions and to query various attributes of the active ActivitySession associated with the thread.
3. Set the timeout, in seconds, after which any subsequently started ActivitySessions are automatically completed by the ActivitySession service. If the session bean or J2EE client does not specifically set this value, the default timeout (300 seconds) is used.

The default timeout can also be overridden for each application server, on the **server-> Activity Session Service** panel of the administrative console.

4. Start the ActivitySession, by calling the `beginSession()` method of the UserActivitySession.
5. Within the ActivitySession, call business methods to do the work needed. You can also call UserActivitySession methods to manage the ActivitySession; for example, to get the status of the ActivitySession or to checkpoint all the ActivitySession resources involved in the ActivitySession.
6. End the ActivitySession, by calling the `endSession()` method of the UserActivitySession.

The following code extract provides a basic example of using the UserActivitySession interface:

```
// Get initial context
InitialContext ic = new InitialContext();
// Lookup UserActivitySession
UserActivitySession uas = (UserActivitySession)ic.lookup("java:comp/websphere/UserActivitySession");

// Set the ActivitySession timeout to 60 seconds
uas.setSessionTimeout(60);
// Start a new ActivitySession context
uas.beginSession();
// Do some work under this context
MyBeanA beanA.doSomething();
...
MyBeanB beanB.doSomethingElse();
// End the context
uas.endSession(EndModeCheckpoint);
```

---

## Setting EJB module ActivitySession deployment attributes with the Assembly Toolkit

Use this task to set the ActivitySession deployment attributes for an enterprise bean to enable the bean to participate in an ActivitySession context and support ActivitySession-based operations.

You can configure the deployment attributes of an application by using the Deployment Descriptor Editor of WebSphere Studio Application Developer or the Assembly Toolkit.

This topic describes the use of the Assembly Toolkit to configure the ActivitySession deployment attributes. These attributes are in addition to other deployment attributes, like Load at (which specifies when the bean loads its state from the database). This task description assumes that you have an EAR file, which contains an application enterprise bean that can be deployed in WebSphere Application Server. For more

details about using the Assembly Toolkit, see *Assembling applications with the Assembly Toolkit*. For more detail about the fields in the Assembly Toolkit, and for associated task help, see the help information provided with the toolkit.

To set the `ActivitySession` deployment attributes for an enterprise bean, use the Assembly Toolkit to complete the following steps:

1. Start the Assembly Toolkit.
2. Create or edit the application EAR file.

**Note:** Ensure that you set the target server as WBI Server Foundation v5.1.

For example, to change attributes of an existing application, use the import wizard to import the EAR file into the Assembly Toolkit. To start the import wizard:

- a. Click **File-> Import-> EAR file**
  - b. Click **Next**, then select the EAR file.
  - c. In the Target server field, select WBI Server Foundation v5.1
  - d. Click **Finish**
3. In the J2EE Hierarchy view of the J2EE perspective, right-click the EJB module for the enterprise bean instance, then click **Open With > Deployment Descriptor Editor**. A property dialog notebook for the enterprise bean instance is displayed in the property pane.
  4. In the property pane, select the Beans tab.
  5. Select the bean that you want to change.
  6. In the WebSphere Extensions section, under **Bean Cache**, set the **Activate at** attribute to **ActivitySession**:
- An enterprise bean with this activation policy is activated and passivated as follows:
- On an `ActivitySession` boundary, if an `ActivitySession` context is present on activation.
  - On a transaction boundary, if a transaction context, but no `ActivitySession` context, is present on activation.
  - Otherwise on an invocation boundary.
7. In the Local Transactions group box, set the **Boundary** attribute to **ActivitySession**: When this setting is used, the local transaction must be resolved within the scope of any `ActivitySession` in which it was started or, if no `ActivitySession` context is present, within the same bean method in which it was started.
  8. For entity beans, or session beans, set the `ActivitySessions` properties for each EJB method.

- a. In the property pane, select the `ActivitySession` tab.
- b. In the **Configure `ActivitySession` policies** field, click **Add** or **Edit** to set the **`ActivitySession` kind** attribute for methods of the enterprise bean. This specifies how the container must manage the `ActivitySession` boundaries when delegating a method invocation to an enterprise bean's business method:

**Never** The container invokes bean methods without an `ActivitySession` context.

- If the client invokes a bean method from within an `ActivitySession` context, the container throws an `InvalidActivityException` exception, which is a `javax.rmi.RemoteException`.
- If the client invokes a bean method from outside an `ActivitySession` context, the container behaves in the same way as if the **Not Supported** value was set. The client must call the method without an `ActivitySession` context.

#### **Mandatory**

The container always invokes the bean method within the `ActivitySession` context associated with the client. If the client attempts to invoke the bean method without an `ActivitySession` context, the container throws an `ActivityRequiredException` exception to the client. The `ActivitySession` context is passed to any EJB object or resource accessed by an enterprise bean method.

The `ActivityRequiredException` exception is `javax.rmi.RemoteException`.

#### **Requires new**

The container always invokes the bean method within a new `ActivitySession` context, regardless of whether the client invokes the method within or outside an `ActivitySession` context. The new `ActivitySession` context is passed to any enterprise bean objects or resources that are used by this bean method.

Any received `ActivitySession` context is suspended for the duration of the method and resumed after the method ends. The container starts a new `ActivitySession` before method dispatch and completes it after the method ends.

#### **Required**

The container invokes the bean method within an `ActivitySession` context. If a client invokes a bean method from within an `ActivitySession` context, the container invokes the bean method within the client `ActivitySession` context. If a client invokes a bean method outside an `ActivitySession` context, the container creates a new `ActivitySession` context and invokes the bean method from within that context. The `ActivitySession` context is passed to any enterprise bean objects or resources that are used by this bean method.

#### **Not supported**

The container invokes bean methods without an `ActivitySession` context. If a client invokes a bean method from within an `ActivitySession` context, the container suspends the association between the `ActivitySession` and the current thread before invoking the method on the enterprise bean instance. The container then resumes the suspended association when the method invocation returns. The suspended `ActivitySession` context is not passed to any enterprise bean objects or resources that are used by this bean method.

#### **Supports**

If the client invokes the bean method within an `ActivitySession`, the container invokes the bean method within an `ActivitySession` context. If the client invokes the bean method without a `ActivitySession` context, the container invokes the bean method without an `ActivitySession` context. The `ActivitySession` context is passed to any enterprise bean objects or resources that are used by this bean method.

- c. Click **Next**.
- d. Select the methods to which the `ActivitySession` kind policy is to be applied.
- e. Click **Finish**.
- f.

How the container manages the `ActivitySession` boundaries when delegating a method invocation depends on both the **ActivitySession kind** attribute, set here, and the **Transaction attribute** attribute, as described in [Setting transactional attributes in the deployment descriptor](#). For more detail about the relationship between these two properties, see [Combining transaction and ActivitySession container policies](#).

9. Save your changes to the deployment descriptor.
  - a. Close the Deployment Descriptor Editor.
  - b. When prompted, click **Yes** to indicate that you want to save changes to the deployment descriptor.
10. Verify the archive files.
11. Generate code for deployment for EJB modules or for enterprise applications that use EJB modules.
12. Optional: Test your completed module on a WebSphere Application Server installation. Right-click a module, click **Run on Server**, and follow the instructions in the displayed wizard. Note that **Run on Server** works on the Windows, Linux/Intel, and AIX operating systems only; you cannot deploy remotely from the Assembly Toolkit to a WebSphere Application Server installation on a UNIX operating system such as Solaris.

#### **Important**

**Important:** Use **Run On Server** for unit testing only. Assembly Server Toolkit controls the WebSphere Application Server installation and, when an application is published remotely, the Toolkit overwrites the server configuration file for that server. Do not use on production servers.

For instructions on remote testing, see the article “Setting Up a Remote WebSphere Application Server in WebSphere Studio V5” at [http://www7b.boulder.ibm.com/wsdd/techjournal/0303\\_yuen/yuen.html](http://www7b.boulder.ibm.com/wsdd/techjournal/0303_yuen/yuen.html).

After assembling your application, use a systems management tool to deploy the EAR file onto the application server that is to run the application; for example, using the administrative console as described in Deploying and managing applications.

---

## Setting Web module ActivitySession deployment attributes with the Assembly Toolkit

Use this task to set the ActivitySession deployment attributes for a Web application to start UserActivitySessions and perform work scoped within ActivitySessions.

You can configure the deployment attributes of an application by using the Deployment Descriptor Editor of WebSphere Studio Application Developer or the Assembly Toolkit.

This topic describes the use of the Assembly Toolkit to configure the deployment attributes. This task description assumes that you have an EAR file, which contains an application enterprise bean that can be deployed in WebSphere Application Server. For more details about using the Assembly Toolkit, see Assembling applications with the Assembly Toolkit.

To set the ActivitySession deployment attributes for a Web application, use the Assembly Toolkit to complete the following steps:

1. Start the Assembly Toolkit.
2. Create or edit the Web module. For example, to change attributes of an existing module, click **File-> Open** then select the archive file for the module. For example, to change attributes of an existing module, use the import wizard to import the EAR or WAR file into the Assembly Toolkit. To start the import wizard:
  - a. Click **File-> Import-> WAR file**
  - b. Click **Next**, then select the WAR file.
  - c. Click **Finish**
3. In the J2EE Hierarchy view, right-click the Web module, then click **Open With > Deployment Descriptor Editor**. A property dialog notebook for the Web module is displayed in the property pane.
4. In the property pane, select the Extended services tab.
5. Select the servlet that you want to change.
6. In the ActivitySession section, set the **ActivitySession control kind** attribute to either Application, Container, or None.

### Application

The Web application is responsible for starting and ending ActivitySessions, as follows:

- If an HttpSession is active when an application begins an ActivitySession, then the container associates the ActivitySession with the HttpSession.
- If an ActivitySession is started in the absence of an HttpSession, then the application must ensure it is completed before the dispatched method completes; otherwise, an exception results.
- If an HttpSession is associated with a request dispatched to an application with this ActivitySession control value, and if that HttpSession has an ActivitySession associated

with it, then the container dispatches the request in the context of that `ActivitySession`. For example, the container resumes the `ActivitySession` context onto the thread before the dispatch.

- A Web application can use both transactions and `ActivitySessions`. Any transactions started within the scope of an `ActivitySession` must be ended by the web component that started them and within the same request dispatch.

#### **Container**

A servlet has no access to `UserActivitySessions`. Any `HttpSession` started by the servlet has an `ActivitySession` automatically associated with it by the container, and this `ActivitySession` is put onto the thread of execution. If such a servlet is dispatched by a request that has an `HttpSession` containing no `ActivitySession`, then the container starts an `ActivitySession` and associates it with the `HttpSession` and the thread.

A Web application can use both transactions and `ActivitySessions`. Any transactions started within the scope of an `ActivitySession` must be ended by the web component that started them and within the same request dispatch.

**None** A servlet has no access to `UserActivitySession`. An `HttpSession` started by the servlet does not have an `ActivitySession` automatically associated with it by the container. If such a servlet is dispatched by a request that has an `HttpSession` containing an `ActivitySession`, then the container dispatches the request in the context of that `ActivitySession`. For example, the container resumes the `ActivitySession` context onto the thread before the dispatch.

7. To apply the changes and close the Assembly Toolkit, click **OK**. Otherwise, to apply the values but keep the property dialog open for additional edits, click **Apply**.
8. Save your changes to the deployment descriptor.
  - a. Close the deployment descriptor editor.
  - b. When prompted, click **Yes** to indicate that you want to save changes to the deployment descriptor.
9. Verify the archive files.
10. Generate code for deployment for EJB modules or for enterprise applications that use EJB modules.
11. Optional: Test your completed module on a WebSphere Application Server installation. Right-click a module, click **Run on Server**, and follow the instructions in the displayed wizard. Note that **Run on Server** works on the Windows, Linux/Intel, and AIX operating systems only; you cannot deploy remotely from the Assembly Toolkit to a WebSphere Application Server installation on a UNIX operating system such as Solaris.

#### **Important**

**Important:** Use **Run On Server** for unit testing only. Assembly Server Toolkit controls the WebSphere Application Server installation and, when an application is published remotely, the Toolkit overwrites the server configuration file for that server. Do not use on production servers.

For instructions on remote testing, see the article “Setting Up a Remote WebSphere Application Server in WebSphere Studio V5” at

[http://www7b.boulder.ibm.com/wsdd/techjournal/0303\\_yuen/yuen.html](http://www7b.boulder.ibm.com/wsdd/techjournal/0303_yuen/yuen.html).

After assembling your application, use a systems management tool to deploy the WAR file; for example, using the administrative console as described in *Deploying and managing applications*.

---

## **Disabling or enabling the ActivitySession service**

Use this task to disable or enable the `ActivitySession` service for an application server.

You can use the `ActivitySession Startup` property to specify whether or not the `ActivitySession` service is started automatically for an application server.



To configure the ActivitySession **Startup** property for an application server, use the Administrative console to complete the following steps:

1. Start the Administrative console.
2. In the navigation pane, expand **Servers-> Manage Application Servers** This displays a list of the application servers in the content pane.
3. In the Content pane, select the application server that you want to configure. This displays the properties for the application server in the content pane.
4. In the Additional Properties table, select **ActivitySession service**. This displays the ActivitySession service properties in the content pane.
5. Select or clear the **Startup** property as needed:

**Selected**

[Default] The ActivitySession service is started when the application server is started. This enables applications that specify use of ActivitySessions in their deployment descriptors to run on such an application server.

**Cleared**

The ActivitySession service is not started when the application server is started. Applications that specify use of ActivitySessions in their deployment descriptors cannot start on such an application server.

Any attempt to start an application that uses ActivitySessions is rejected and a message issued:

```
WACS0043E: Error found starting an application.
application_name specified an ActivitySession attribute that is not allowed
when the ActivitySession service is not enabled
```

If this happens during server startup, the server continues to start without the application.

6. Click **OK**.
7. To save your configuration, click **Save** on the task bar of the Administrative console window.
8. To have the changed configuration take effect, stop then restart the application server.

---

## Configuring the default ActivitySession timeout for an application server

Use this task to configure the default ActivitySession timeout for an application server, after which any started ActivitySessions are completed automatically by the ActivitySession service.

The ActivitySession timeout is used to reset any ActivitySession whose remote client has failed to complete the ActivitySession in a timely fashion. The initial default timeout can be configured separately for each application server, and can be overridden programmatically by the UserActivitySession interface (setSessionTimeout). If an ActivitySession that contains a transaction reaches the timeout, the transaction's timeout is accelerated so that it is timed out (and rolled back) immediately before the ActivitySession is reset.

To configure the default ActivitySession timeout for an application server, use the WebSphere Administrative console to complete the following steps:

1. Start the WebSphere Administrative console.
2. In the navigation pane, expand **Servers-> Manage Application Servers** This displays a list of the application servers in the content pane.
3. In the Content pane, select the application server that you want to configure. This displays the properties for the application server in the content pane.
4. In the Additional Properties table, select **ActivitySession service**. This displays the ActivitySession service properties in the content pane.
5. Set the **ActivitySession timeout** property to the default timeout as an integer number of seconds.



- -1 indicates that ActivitySessions never timeout
- 0 indicates that the default timeout, 300 seconds, applies
- Other values are an integer number of seconds

6. Click **OK**.
7. To save your configuration, click **Save** on the task bar of the Administrative console window.
8. To have the changed configuration take effect, stop then restart the application server.

## ActivitySession service settings

Use this page to administer the run-time properties of the ActivitySession service.

To view this administrative console page, click **Servers > Application Servers > server\_name > Activity Session Service**.

### Startup

Specifies whether the server will attempt to start the ActivitySession service.

<b>Default</b>	Selected
<b>Range</b>	<p><b>Selected</b></p> <p>When the application server starts, it attempts to start the ActivitySession service automatically.</p> <p><b>Cleared</b></p> <p>The server does not try to start the ActivitySession service. If ActivitySessions are to be used in applications that run on this server, the system administrator must start the service manually or select this property then restart the server.</p>

### Default timeout

The default timeout for an ActivitySession. A server resets an ActivitySession if a remote client has failed to complete the ActivitySession within this time period.

The Default ActivitySession timeout specifies the time after which an ActivitySession is completed automatically by the ActivitySession service, if a remote client has failed to complete the ActivitySession within the specified time. The initial default timeout can be configured separately for each application server, and can be overridden programmatically by the UserActivitySession interface (setSessionTimeout).

<b>Data type</b>	Integer
<b>Units</b>	Seconds
<b>Default</b>	300 (5 minutes)
<b>Range</b>	<p>-1 through 2147483647 seconds</p> <ul style="list-style-type: none"> <li>• -1 indicates that ActivitySessions never timeout</li> <li>• 0 indicates that the default timeout applies</li> <li>• Other values are an integer number of seconds</li> </ul>

---

## Troubleshooting ActivitySessions

Use this overview task to help resolve a problem that you think is related to the ActivitySession service.

To identify and resolve ActivitySession-related problems, you can use the standard WebSphere Application Server RAS facilities. If you encounter a problem that you think might be related to ActivitySessions, complete the following stages:

1. Check for ActivitySession messages in the admin console. The ActivitySession service produces diagnostic messages prefixed by “WACS”. The error message indicates the nature of the problem and provides some detail. The associated message information provides an explanation and any user actions to resolve the problem.
2. Check for ActivitySession messages in the activity log. Activity log messages produced by the ActivitySession service are accompanied by Log Analyzer descriptions.
3. Check for more messages in the application server’s stdout.log. For more information about a problem, check the stdout.log file for the application server, which should contain more error messages and extra details about the problem.

---

## The ActivitySession service application programming interfaces

The ActivitySession service consists of an application programming interface available to Web applications, session EJBs, and J2EE client applications for application-managed demarcation of ActivitySession context.

Applications use the UserActivitySession interface, which provides demarcation scope methods.

### ActivitySession API

The ActivitySession service provides the UserActivitySession interface for use by EJB Session beans using bean-managed context demarcation, Web application components configured with **ActivitySession control=Web Application**, and J2EE client applications. This UserActivitySession interface defines the set of ActivitySession operations available to an application component. An implementation of this interface is obtained via a JNDI lookup of the URL “java:comp/websphere/UserActivitySession”. It is used to begin and end ActivitySessions and to query various attributes of the active ActivitySession associated with the thread.

For more information about the ActivitySession API, see the ActivitySession API (Javadoc).

The ActivitySession API and the implementation of its interfaces is contained in the com.ibm.websphere.ActivitySession package.

### Programming Examples

The following code extract provides a basic example of using the UserActivitySession interface:

```
// Get initial context
InitialContext ic = new InitialContext();
// Lookup UserActivitySession
UserActivitySession uas = (UserActivitySession)ic.lookup("java:comp/websphere/UserActivitySession");

// Set the ActivitySession timeout to 60 seconds
uas.setSessionTimeout(60);
// Start a new ActivitySession context
uas.beginSession();
// Do some work under this context
MyBeanA beanA.doSomething();
...
MyBeanB beanB.doSomethingElse();
// End the context
uas.endSession(EndModeCheckpoint);
```

---

## Samples: ActivitySessions

This topic describes the ActivitySession samples provided with WebSphere Application Server.

### MasterMind sample

This sample is based on the game MasterMind. It consists of the following components:

- A servlet, configured with Activity session control set to Container, that accesses a stateful session bean.
- A stateful session bean, configured with an activation policy of ActivitySession containing transient state data.

The servlet begins an HttpSession at the start of each new game, and ends it at the end of each game; therefore an ActivitySession lasts for the duration of each game. The ActivitySession activation policy stops the bean from being passivated and therefore the transient data remains in memory. This is to demonstrate HttpSession/ActivationSession association in the web container, and an ActivitySession-scoped activation policy.

#### **J2EE client container application and a CMP entity bean backed by a one-phase commit datasource**

In this sample, the entity bean is configured with the following properties:

- TX\_NOT\_SUPPORTED
- An ActivitySession container managed policy of REQUIRES
- An LTC boundary of ActivitySession
- An LTC Resolution Control of ContainerAtBoundary

The client accesses the UserActivitySession, begins an ActivitySession, updates two instances of the bean, then ends the ActivitySession. It does this twice using EndModeReset then EndModeCheckpoint. This sample demonstrates the following functionality:

- Client access to the UserActivitySession interface
- Multiple RMLTs being scoped to the ActivitySession and automatically taking their completion direction from that of the ActivitySession

The entity bean also holds a transient variable incremented by each method call (gets and sets for the persistent data). This value is checked before the end of the ActivitySession to show that the same bean instance is used. The client checks for the correct results.

#### **A J2EE client container application and two session beans with different ActivitySession types**

This sample consists of a J2EE client container application and the following session beans:

- SLB1, a stateless session bean configured with an ActivitySession Type of Bean.
- SFB2, a stateful session bean configured with ActivitySession Type of Requires, an LTC boundary of ActivitySession, LTC Resolution Control of APPLICATION, and an LTC Unresolved Action of ROLLBACK.

Both beans are configured with TX\_NOTSUPPORTED.

This sample performs the following steps:

1. The client starts SLB1
2. SLB1 accesses the UserActivitySession interface, begins an ActivitySession, then calls a method on SFB2
3. SFB2 accesses the UserActivitySession interface, begins an ActivitySession, calls a method on SFB2
4. SFB2 gets a connection (setAutoCommit false) then uses JDBC to update a single-phase datasource.
5. SLB1 then optionally calls a separate method on SFB2 to finish the work either committing or rolling-back the RMLT.
6. SLB1 then ends the ActivitySession with an EndModeCheckpoint.

This demonstrates that the ActivitySession completion direction is unconnected to the direction of the RMLTs, although their containment is bound to the ActivitySession, and the use of the container using the unresolved action when the RMLT is not completed. It also shows a bean-managed ActivitySessions bean using the UserActivitySession interface. The sample checks for correct results and reports them back to the client.

---

## ActivitySession service: Resources for learning

Use the links in this topic to find relevant supplemental information about ActivitySessions. The information resides on IBM and non-IBM Internet sites, whose sponsors control the technical accuracy of the information.

These links are provided for convenience. Often, the information is not specific to the IBM WebSphere Application Server product, but is useful all or in part for understanding the product. When possible, links are provided to technical papers and Redbooks that supplement the broad coverage of the release documentation with in-depth examinations of particular product areas.

View links to additional information about:

- “Programming model and decisions”
- “Programming specifications”
- “Other”

### Programming model and decisions

- ActivitySession API (Javadoc)

### Programming specifications

- J2EE Activity Service for Extended Transactions
- Java Transaction API (JTA) 1.0.1

### Other

- WebSphere Application Server Enterprise Version 5 Overview: Advanced Transactional Connectivity
- Listing of PDF files to learn about WebSphere Application Server Version 5
- Listing of all IBM WebSphere Application Server Redbooks
- Listing of all IBM WebSphere Application Server Whitepapers
- WebSphere Application Server Enterprise Edition 4.0: A Programmer’s Guide



---

## Chapter 26. Using naming

Naming is used by clients of WebSphere Application Server applications most commonly to obtain references to objects related to those applications, such as Enterprise JavaBeans (EJB) homes. The following steps outline the context of Naming in the overall application development and deployment process. Steps for this task follow:

1. Develop your application using either JNDI or CosNaming (CORBA) interfaces.  
Use these interfaces to look up server application objects that are bound into the name space and obtain references to them. Most Java developers use the JNDI interface. However, the CORBA CosNaming interface is also available for performing Naming operations on WebSphere Application Server name servers or other CosNaming name servers.
2. Assemble your application using the Assembly Toolkit. Application assembly is a packaging and configuration step that is a prerequisite to application deployment. If the application you are assembling is a client to an application running in another process, you should qualify the `jndiName` values in the deployment descriptors for the objects related to the other application. Otherwise, you may need to override the names with qualified names during application deployment. If the objects have fixed qualified names configured for them, you should use them so that the `jndiName` values do not depend on the other application's location within the topology of the cell.
3. Deploy your application  
Put your assembled application onto the application server. If the application you are assembling is a client to an application running in another server process, be sure to qualify the `jndiName` values for the other application's server objects if they are not already qualified.  
For more information on qualified names, see "Lookup names support in deployment descriptors and thin clients."
4. Configure name space bindings. This step is necessary in these cases:
  - Your deployed application is to be accessed by legacy client applications running on previous versions of WebSphere Application Server. In this case, you must configure additional name bindings for application objects relative to the default initial context for legacy clients. (Version 5 clients have a different initial context from legacy clients.)
  - The application requires qualified name bindings for such reasons as:
    - It will be accessed by J2EE client applications or server applications running in another server process.
    - It will be accessed by thin client applications.In this case, you can configure name bindings as additional bindings for application objects. The qualified names for the configured bindings are *fixed*, meaning they do not contain elements of the cell topology that can change if the application is moved to another server. Objects as bound into the name space by the system can always be qualified with a topology-based name. You must explicitly configure a name binding to use as a fixed qualified name.  
For more information on qualified names, see "Lookup names support in deployment descriptors and thin clients." For more information on configured name bindings, see "Configured name bindings."
5. Troubleshoot any problems that develop.  
If a Naming operation is failing and you need to verify whether certain name bindings exist, use the `dumpNameSpace` tool to generate a dump of the name space.

---

### Naming

Naming is used by clients of WebSphere Application Server applications to obtain references to objects related to those applications, such as Enterprise JavaBeans (EJB) homes.

These objects are bound into a mostly hierarchical structure, referred to as a *name space*. In this structure, all non-leaf objects are called *contexts*. Leaf objects can be contexts and other types of objects.

Naming operations, such as lookups and binds, are performed on contexts. All naming operations begin with obtaining an *initial context*. You can view the initial context as a starting point in the name space.

The name space structure consists of a set of *name bindings*, each consisting of a name relative to a specific context and the object bound with that name. For example, the name `myApp/myEJB` consists of one non-leaf binding with the name `myApp`, which is a context. The name also includes one leaf binding with the name `myEJB`, relative to `myApp`. The object bound with the name `myEJB` in this example happens to be an EJB home reference. The whole name `myApp/myEJB` is relative to the initial context, which you can view as a starting place when performing naming operations.

You can access and manipulate the name space through a *name server*. Users of a name server are referred to as *naming clients*. Naming clients typically use the Java Naming and Directory Interface (JNDI) to perform naming operations. Naming clients can also use the Common Object Request Broker Architecture (CORBA) CosNaming interface.

Typically, objects bound to the name space are resources and objects associated with installed applications. These objects are bound by the system, and client applications perform lookup operations to obtain references to them. Occasionally, server and client applications bind objects to the name space. An application can bind objects to transient or persistent partitions, depending on requirements.

In J2EE environments, some JNDI operations are performed with `java:` URL names. Names bound under these names are bound to a completely different name space which is local to the calling process. However, some lookups on the `java:` name space may trigger indirect lookups to the name server.

---

## Version 5 features for name space support

The following are features of the WebSphere Application Server new to naming implementation as of Version 5:

- **Name space is distributed.**

For additional scalability, the name space for a cell is distributed among various servers. Every server has a name server. In previous releases, there was only one name server for an entire administrative domain.

In WebSphere Application Server versions prior to V5, all servers shared the same default initial context, and everything was bound relative to that same initial context. In WebSphere Application Server V5, the default initial context for a server is its server root. System artifacts, such as EJB homes and resources, are bound to the server root of the server with which they are associated.

- **Transient and persistent partitions.**

The name space is partitioned into transient areas and persistent areas. Server roots are transient. System-bound artifacts such as EJB homes and resources are bound under server roots. There is a cell persistent root, which you can use for cell-scoped persistent bindings, and a node persistent root, which you can use to bind objects with a node scope.

- **System name space structure.**

The name space for the entire cell is federated among all servers in the cell. Every server process contains a name server. All name servers provide the same logical view of the cell name space. The various server roots and persistent partitions of the name space are interconnected by means of a system name space. You can use the system name space structure to traverse to any context in the cell name space.

- **Configured bindings.**

You can use the configuration graphical interface and script interfaces to configure bindings in various root contexts within the name space. These bindings are read-only and are bound by the system at server startup.

- **Support for CORBA Interoperable Naming Service (INS) object URLs.**



WebSphere Application Server V5 contains support for Common Object Request Broker Architecture (CORBA) object URLs (corbaloc and corbname) as Java Naming and Directory Interface (JNDI) provider URLs and lookup names.

## Name space logical view

The name space for the entire cell is federated among all servers in the cell. Every server process contains a name server. All name servers provide the same logical view of the cell name space. The various server roots and persistent partitions of the name space are interconnected by a system name space. You can use the system name space structure to traverse to any context in a the cell's name space. A logical view of the name space is shown in the following diagram.

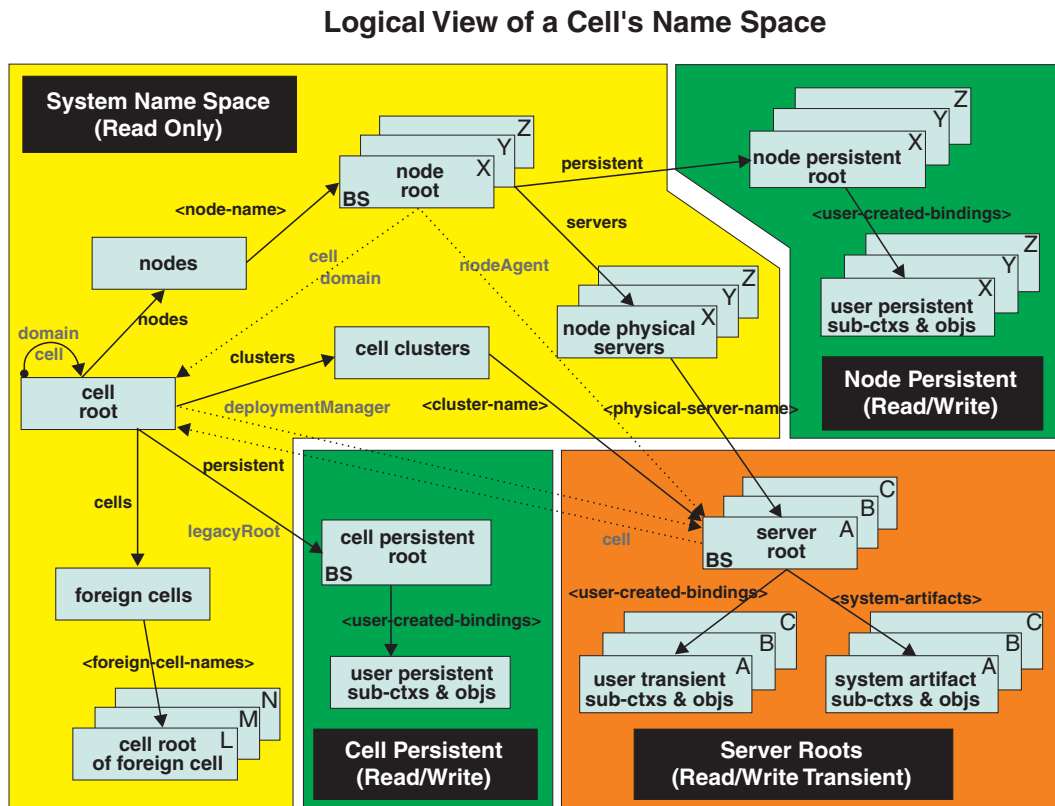


Figure 34. Name Space Logical View

The bindings in the preceding diagram appear with solid arrows, labeled in bold, and dashed arrows, labeled in gray. Solid arrows represent *primary bindings*. A primary binding is formed when the associated subcontext is created. Dashed arrows show *linked bindings*. A linked binding is formed when an existing context is bound under an additional name. Linked bindings are added for convenience or interoperability with previous WebSphere Application Server versions.

A cell name space is composed of contexts which reside in servers throughout the cell. All name servers in the cell provide the same logical view of the cell name space. A name server constructs this view at startup by reading configuration information. Each name server has its own local in-memory copy of the name space and does not require another running server to function. There are, however, a few exceptions. Server roots for other servers are not replicated among all the servers. The respective server for a server root must be running to access that server root context.

In WebSphere Application Server Network Deployment cells, the cell and node persistent areas can be read even if the deployment manager and respective node agent are not running. However, the

deployment manager must be running to update the cell persistent segment, and a node agent must be running to update its respective node persistent segment.

## **Name space partitions**

There are four major partitions in a cell name space:

- System name space partition
- Server roots partition
- Cell persistent partition
- Node persistent partition

### **System name space partition**

The system name space contains a structure of contexts based on the cell topology. The system structure supports traversal to all parts of a cell name space and to the cell root of other cells, which are configured as foreign cells. The root of this structure is the cell root. In addition to the cell root, the system structure contains a node root for each node in the cell. You can access other contexts of interest specific to a node from the node root, such as the node persistent root and server roots for servers configured in that node.

All contexts in the system name space are read-only. You cannot add, update, or remove any bindings.

### **Server roots partition**

Each server in a cell has a server root context. A server root is specific to a particular server. You can view the server roots for all servers in a cell as being in a transient read/write partition of the cell name space. System artifacts, such as EJB homes for server applications and resources, are bound under the server root context of the associated server. A server application can also add bindings under its server root. These bindings are transient. Therefore, the server application creates all required bindings at application startup, so they exist anytime the application is running.

A server cluster is composed of many servers that are logically equivalent. Each member of the cluster has its own server root. These server roots are not replicated across the cluster. In other words, adding a binding to the server root of one member does not propagate it to the server roots of the other cluster members. To maintain the same view across the cluster, you should create all user bindings under the server root by the server application at application startup so that the bindings are present under the server root of each cluster member. Because of Workload Management (WLM) behavior, a JNDI client outside a cluster has no control over which cluster member's server root context becomes the target of the JNDI operation. Therefore, you should execute bind operations to the server root of a cluster member from within that cluster member process only.

Distributing application objects among many server roots is a departure from previous WebSphere Application Server releases, where all system artifacts were bound under a single root. This change can affect the names that clients use to look up these objects.

Server-scoped bindings are relative to a server's server root.

### **Cell persistent partition**

The root context of the cell persistent partition is the cell persistent root. A binding created under the cell persistent root is saved as part of the cell configuration and continues to exist until it is explicitly removed. Applications that need to create additional persistent bindings of objects generally associated with the cell can bind these objects under the cell persistent root.

It is important to note that the cell persistent area is not designed for transient, rapidly changing bindings. The bindings are more static in nature, such as part of an application setup or configuration, and are not created at run time.

An important role of the cell persistent root is as the initial context for clients running in previous WebSphere Application Server versions. If you want to access an enterprise bean by WebSphere

Application Server v4.0.x and 3.5.x clients, you must ensure that a binding for it has been added to the cell persistent root. You can configure these additional bindings as cell-scoped bindings.

### Node persistent partition

The node persistent partition is similar to the cell partition except that each node has its own node persistent root. A binding created under a node persistent root is saved as part of that node configuration and continues to exist until it is explicitly removed.

Applications that need to create additional persistent bindings of objects associated with a specific node can bind those objects under that particular node's node persistent root. As with the cell persistent area, it is important to note that the node persistent area is not designed for transient, rapidly changing bindings. These bindings are more static in nature, such as part of an application setup or configuration, and are not created at run time.

Unlike the cell persistent root, the node persistent root plays no special role in interoperability with WebSphere Application Server clients of previous releases. Node-scoped bindings are relative to a node's node persistent root.

---

## Initial context support

All naming operations begin with obtaining an initial context. You can view the initial context as a starting point in the name space. Use the initial context to perform naming operations, such as looking up and binding objects in the name space.

### Initial contexts registered with the ORB as initial references

The server root, cell persistent root, cell root, and node root are registered with the name server's ORB and can be used as an initial context. An initial context is used by CORBA and enterprise bean applications as a starting point for name space lookups. The keys for these roots as recognized by the ORB are shown in the following table:

Root Context	Initial Reference Key
Server Root	NameServiceServerRoot
Cell Persistent Root	NameServiceCellPersistentRoot
Cell Root	NameServiceCellRoot, NameService
Node Root	NameServiceNodeRoot

A server root initial context is the server root context for the specific server you are accessing. Similarly, a node root initial context is the node root for the server being accessed.

You can use the previously mentioned keys in CORBA INS object URLs (corbaloc and corbaname) and as an argument to an ORB `resolve_initial_references` call. For examples, see CORBA and JNDI programming examples, which show how to get an initial context.

### Default initial contexts

The default initial context depends on the type of client. Different categories of clients and the corresponding default initial context follow.

- **WebSphere Application Server V5 JNDI interface implementation**

The JNDI interface is used by EJB applications to perform name space lookups. WebSphere Application Server clients by default use the WebSphere Application Server CosNaming JNDI plug-in implementation. The default initial context for clients of this type is the server root of the server specified by the provider URL. For more details, refer to the JNDI programming examples on getting initial contexts.

- **WebSphere Application Server JNDI interface implementation prior to V5**

WebSphere Application Server clients running in releases prior to WebSphere Application Server V5 by default use WebSphere Application Server's v4.0 CosNaming JNDI plug-in implementation. The default initial context for clients of this type is the cell persistent root, also known as the *legacy root*.

- **Other JNDI implementation**

Some applications can perform name space lookups with a non-WebSphere Application Server CosNaming JNDI plug-in implementation. Assuming the key **NamingContext** is used to obtain the initial context, the default initial context for clients of this type is the cell root.

- **CORBA**

The standard CORBA client obtains an initial `org.omg.CosNaming.NamingContext` reference with the key **NamingContext**. The initial context in this case is the cell root.

---

## Lookup names support in deployment descriptors and thin clients

Server objects, such as EJB homes, are bound relative to the server root context for the server in which the application is installed. Other objects, such as resources, can also be bound to a specific server root. The names used to look up these objects must be qualified so as to select the correct server root. This is a departure from previous versions of WebSphere Application Server, where these objects were all bound under a single root context. This section discusses what relative and qualified names are, when they can be used, and how you can construct them.

### Relative names

All names are relative to a context. Therefore, a name that can be resolved from one context in the name space cannot necessarily be resolved from another context in the name space. This point is significant because the system binds objects with names relative to the server root context of the server in which the application is installed. Each server has its own server root context. The initial JNDI context is by default the server root context for the server identified by the provider URL used to obtain the initial context. (Typically, the URL consists of a host and port.) For applications running in a server process, the default initial JNDI context is the server root for that server. A relative name will resolve successfully when the initial context is obtained from the server which contains the target object, but it will not resolve successfully from an initial context obtained from another server.

If all clients of a server application run in the same server process as the application, all objects associated with that application are bound to the same initial context as the clients' initial context. In this case, only names relative to the server's server root context are required to access these server objects. Frequently, however, a server application has clients that run outside the application's server process. The initial context for these clients can be different from the server application's initial context, and lookups on the relative names for server objects may fail. These clients need to use the qualified name for the server objects. This point must be considered when setting up the `jndiName` values in a J2EE client application deployment descriptors and when constructing lookup names in thin clients. Qualified names resolve successfully from any initial context in the cell.

### Qualified names

All names are relative to a context. Here, the term *qualified name* refers to names that can be resolved from any initial context in a cell. This action is accomplished by using names that navigate to the same context, the cell root. The rest of the qualified name is then relative to the cell root and uniquely identifies an object throughout the cell. All initial contexts in a server (that is, all naming contexts in a server registered with the ORB as an initial reference) contain a binding with the name **cell**, which links back to the cell root context. All qualified names begin with the string **cell/** to navigate from the current initial context back to the cell root context.

A qualified name for an object is the same throughout the cell. The name can be topology-based, or some fixed name bound under the cell persistent root. Topology-based names, described in more detail below, navigate through the system name space to reach the target object. A fixed name bound under the cell

persistent root has the same qualified name throughout the cell and is independent of the topology. Creating a fixed name under the cell persistent root for a server application object requires an extra step when the server application is installed, but this step eliminates impacts to clients when the application is moved to a different location in the cell topology. The process for creating a fixed name is described later in this section.

Generally speaking, you **must** use qualified names for EJB `jndiName` values in a J2EE client application deployment descriptors and for EJB lookup names in thin clients. The only exception is when the initial context is obtained from the server in which the target object resides. For example, a session bean which is a client to an entity bean can use a relative name if the two beans run in the same server. If the session bean and entity beans run in different servers, the `jndiName` for the entity bean must be qualified in the session bean's deployment descriptors. The same requirement may be true for resources as well, depending on the scope of the resource.

- **Topology-based names**

The system name space partition in a cell's name space reflects the cell's topology. This structure can be navigated to reach any object bound into the cell's name space. Topology-based qualified names include elements from the topology which reflect the object's location within the cell. For a system-bound object, such as an EJB home, the form for a topology-based qualified name depends on whether the object is bound to a single server or cluster. Both forms are described below.

**Single Server**

An object bound in a single server has a topology-based qualified name of the following form:

```
cell/nodes/nodeName/servers/serverName/relativeJndiName
```

where *nodeName* and *serverName* are the node name and server name for the server where the object is bound, and *relativeJndiName* is the unqualified name of the object; that is, the object's name relative to its server's server root context.

**Server Cluster**

An object bound in a server cluster has a topology-based qualified name of the following form:

```
cell/clusters/clusterName/relativeJndiName
```

where *clusterName* is the name of the server cluster where the object is bound, and *relativeJndiName* is the unqualified name of the object; that is, the object's name relative to a cluster member's server root context.

- **Fixed names**

It is possible to create a fixed name for a server object so that the qualified name is independent of the cell topology. This quality is desirable when clients of the application run in other server processes or as pure clients. Fixed names have the advantage of not changing if the object is moved to another server. The `jndiName` values in deployment descriptors for a J2EE client application can reference the qualified fixed name for a server object regardless of the cell topology on which the client or server application is being installed.

Defining a cell-wide fixed name for a server application object requires an extra step after the server application is installed. That is, a binding for the object must be created under the cell persistent root. A fixed name bound under the cell persistent root can be any name, but all names under the cell persistent root must be unique within the cell because the cell persistent root is global to the entire cell.

A qualified fixed name has the form:

```
cell/persistent/fixedName
```

where *fixedName* is an arbitrary fixed name.

The binding can be created programmatically (for example, using JNDI). However, it is probably more convenient to configure a cell-scoped binding for the server object.

You must keep the programmatic or configured binding up-to-date. Configured EJB bindings are based on the location of the enterprise bean within the cell topology, and moving the EJB application to another single server or to a server cluster, for example, requires the configured binding to be updated. Similar changes affect an EJB home reference programmatically bound so that the fixed name would

need to be rebound with a current reference. However, for J2EE clients, the `jndiName` value for the object, and for thin clients, the lookup name for the object, remains the same. In other words, clients that access objects by fixed names are not affected by changes to the configuration of server applications they access.

---

## JNDI support in WebSphere Application Server

IBM WebSphere Application Server includes a name server to provide shared access to Java components, and an implementation of the `javax.naming` JNDI package which supports user access to the WebSphere Application Server name server through the JNDI naming interface.

WebSphere Application Server does **not** provide implementations for:

- `javax.naming.directory` or
- `javax.naming.ldap` packages

Also, WebSphere Application Server does **not** support interfaces defined in the `javax.naming.event` package.

However, to provide access to LDAP servers, the development kit shipped with WebSphere Application Server supports Sun's implementation of:

- `javax.naming.ldap` and
- `com.sun.jndi.ldap.LdapCtxFactory`

WebSphere Application Server's JNDI implementation is based on version 1.2 of the JNDI interface, and was tested with Version 1.2.1 of Sun's JNDI Service Provider Interface (SPI).

The default behavior of this JNDI implementation is adequate for most users. However, users with specific requirements can control certain aspects of JNDI behavior.

---

## Developing applications that use JNDI

References to EJB homes and other artifacts such as data sources are bound to the WebSphere name space. These objects can be obtained through the JNDI interface. Before you can perform any JNDI operations, you need to get an initial context. You can use the initial context to look up objects bound to the WebSphere name space.

These examples describe how to get an initial context and how to perform lookup operations.

- Getting the default initial context
- Getting an initial context by setting the provider URL property
- Setting the provider URL property to select a different root context as the initial context
- Looking up an EJB home with JNDI
- Looking up a JavaMail session with JNDI

In these examples, the default behavior of features specific to WebSphere's JNDI Context implementation is used.

WebSphere Application Server's JNDI context implementation includes special features. JNDI caching enhances performance of repeated lookup operations on the same objects. Name syntax options offer a choice of a name syntaxes, one optimized for typical JNDI clients, and one optimized for interoperability with CosNaming applications. Most of the time, the default behavior of these features is the preferred behavior. However, sometimes you should modify the behavior for specific situations.

JNDI caching and name syntax options are associated with a `javax.naming.InitialContext` instance. To select options for these features, set properties that are recognized by the WebSphere Application Server's



initial context factory. To set JNDI caching or name syntax properties which will be visible to WebSphere Application Server's initial context factory, follow the following steps.

## 1. Optional: Configure JNDI caches

JNDI caching can greatly increase performance of JNDI lookup operations. By default, JNDI caching is enabled. In most situations, this default is the desired behavior. However, in specific situations, use the other JNDI cache options.

Objects are cached locally as they are looked up. Subsequent lookups on cached objects are resolved locally. However, cache contents can become stale. This situation is not usually a problem, since most objects you look up do not change frequently. If you need to look up objects which change relatively frequently, change your JNDI cache options.

JNDI clients can use several properties to control cache behavior.

You can set properties:

- From the command line by entering the actual string value. For example:

```
java -Dcom.ibm.websphere.naming.jndicache.maxentrylife=1440
```

- In a `jndi.properties` file by creating a file named `jndi.properties` as a text file with the desired properties settings. For example:

```
...
com.ibm.websphere.naming.jndicache.cacheobject=none
...
```

Include the file as the beginning of the classpath, so that the class loader loads your copy of `jndi.properties` before any other copies.

- Within a Java program by using the **PROPS.JNDI\_CACHE\*** Java constants, defined in the ***com.ibm.websphere.naming.PROPS*** file. The constant definitions follow:

```
public static final String JNDI_CACHE_OBJECT =
 "com.ibm.websphere.naming.jndicache.cacheobject";
public static final String JNDI_CACHE_OBJECT_NONE = "none";
public static final String JNDI_CACHE_OBJECT_POPULATED = "populated";
public static final String JNDI_CACHE_OBJECT_CLEARED = "cleared";
public static final String JNDI_CACHE_OBJECT_DEFAULT =
 JNDI_CACHE_OBJECT_POPULATED;
```

```
public static final String JNDI_CACHE_NAME =
 "com.ibm.websphere.naming.jndicache.cachename";
public static final String JNDI_CACHE_NAME_DEFAULT = "providerURL";
```

```
public static final String JNDI_CACHE_MAX_LIFE =
 "com.ibm.websphere.naming.jndicache.maxcachelife";
public static final int JNDI_CACHE_MAX_LIFE_DEFAULT = 0;
```

```
public static final String JNDI_CACHE_MAX_ENTRY_LIFE =
 "com.ibm.websphere.naming.jndicache.maxentrylife";
public static final int JNDI_CACHE_MAX_ENTRY_LIFE_DEFAULT = 0;
```

To use the previous properties in a Java program, add the property setting to a hashtable and pass it to the `InitialContext` constructor as follows:

```
java.util.Hashtable env = new java.util.Hashtable();
...

// Disable caching
env.put(PROPS.JNDI_CACHE_OBJECT, PROPS.JNDI_CACHE_OBJECT_NONE); ...
javax.naming.Context initialContext = new javax.naming.InitialContext(env);
```

## 2. Optional: Specify the name syntax

Most WebSphere applications use JNDI to look up EJB objects and do not need to look up objects bound by CORBA applications. Therefore, the default name syntax used for JNDI names is the most convenient. If your application needs to look up objects bound by CORBA applications, you may need to change your name syntax so that all CORBA CosNaming names can be represented.



JNDI clients can set the name syntax by setting a property. The property setting is applied by the initial context factory when you instantiate a new `java.naming.InitialContext` object. Names specified in JNDI operations on the initial context are parsed according to the specified name syntax.

You can set the property:

- From the command line by entering the actual string value. For example:

```
java -Dcom.ibm.websphere.naming.name.syntax=ins
```

- In a `jndi.properties` file by creating a file named `jndi.properties` as a text file with the desired properties settings. For example:

```
...
com.ibm.websphere.naming.name.syntax=ins
...
```

Include the file as the beginning of the classpath, so that the class loader loads your copy of `jndi.properties` before any other copies.

- Within a Java program by using the `PROPS.NAME_SYNTAX*` Java constants, defined in the `com.ibm.websphere.naming.PROPS` file. The constant definitions follow:

```
public static final String NAME_SYNTAX =
 "com.ibm.websphere.naming.name.syntax";
public static final String NAME_SYNTAX_JNDI = "jndi";
public static final String NAME_SYNTAX_INS = "ins";
```

To use the previous properties in a Java program, add the property setting to a hashtable and pass it to the `InitialContext` constructor as follows:

```
java.util.Hashtable env = new java.util.Hashtable();
...
env.put(PROPS.NAME_SYNTAX, PROPS.NAME_SYNTAX_INS); // Set name syntax to INS
...
javax.naming.Context initialContext = new javax.naming.InitialContext(env);
```

## Example: Getting the default initial context

This example below gets the default initial context. That is, no provider URL is passed to the `javax.naming.InitialContext` constructor. The following section explains the process of determining the address of the bootstrap server to use to obtain the initial context.

```
...
import javax.naming.Context;
import javax.naming.InitialContext;
...
Context initialContext = new InitialContext();
...
```

The default initial context returned depends the runtime environment of the JNDI client. The initial context returned in the various environments are listed below:

- Thin client: The server root context of the server running on the local host at port 2809.
- Pure client:
  - The context specified by the `java.naming.provider.url` property passed to `launchClient` command with the `-CCD` command line parameter. The context usually will be the server root context of the server at the address specified in the URL, although it is possible to construct a `corbaname` or `corbaloc` URL which resolves to some other context.
  - If no provider URL was specified, the server root context of the server running on the host and port specified by the `-CCBootstrapHost` `-CCBootstrapPort` command line parameters. The default host is the local host, and the default port is 2809.
- Server process: The server root context for that process.

Even though no provider URL is explicitly specified in the above example, the `InitialContext` may find a provider URL defined in other places that it searches for property settings.

Users of properties which affect ORB initialization should read the rest of this section for a deeper understanding of exactly how initial contexts are obtained, which has changed from previous releases.

### **Determining which server is used to obtain the initial context**

WebSphere Application Server name servers are CORBA CosNaming name servers, and WebSphere Application Server provides a CosNaming JNDI plug-in implementation for JNDI clients to perform naming operations on WebSphere Application Server name spaces. The WebSphere Application Server CosNaming plug-in implementation is selected through a JNDI property that is passed to the InitialContext constructor. This property is `java.naming.factory.initial`, and it specifies the initial context factory implementation to use to obtain an initial context. The factory returns a `javax.naming.Context` instance, which is part of its implementation.

The WebSphere Application Server initial context factory, `com.ibm.websphere.naming.WsnInitialContextFactory`, is typically used by WebSphere Application Server applications to perform JNDI operations. The WebSphere Application Server run-time environment is set up to use this WebSphere Application Server initial context factory if one is not specified explicitly by the JNDI client. When the initial context factory is invoked, an *initial context* is obtained. The following paragraphs explain how the WebSphere Application Server initial context factory obtains the initial context in client and server environments.

- **Understanding the registration of initial references in server processes**

Every WebSphere Application Server has an ORB used to receive and dispatch invocations on objects running in that server. Services running in the server process can register initial references with the ORB. Each initial reference is registered under a key, which is a string value. An initial reference can be any CORBA object. WebSphere Application Server name servers register several initial contexts as initial references under predefined keys. Each name server initial reference is an instance of the interface `org.omg.CosNaming.NamingContext`.

- **Obtaining initial references in pure client processes**

Pure JNDI clients, that is, JNDI clients which are not running in a WebSphere Application Server process, also have an ORB instance. This client ORB instance can be passed to the InitialContext constructor, but typically the initial context factory creates and initializes the client ORB instance transparently. A client ORB can be initialized with initial references, but the initial references most likely resolve to objects running in some server. The initial context factory does not define any default initial references when it initializes an ORB. If the `resolve_initial_references` method is invoked on the client ORB when no initial references have been configured, the method invocation fails. This condition is typical for pure client processes. To obtain an initial NamingContext reference, the initial context factory must invoke `string_to_object` with an IIOP type CORBA object URL, such as `corbaloc:iiop:myhost:2809`. The URL specifies the address of the server from which to obtain the initial context. The host and port information is extracted from the provider URL passed to the InitialContext constructor. If no provider URL is defined, the WebSphere Application Server initial context factory uses the default provider URL of `corbaloc:iiop:localhost:2809`. The `string_to_object` ORB method resolves the URL and communicates with the target server ORB to obtain the initial reference.

- **Obtaining initial references in server processes**

If the JNDI client is running in a WebSphere Application Server process, the initial context factory obtains a reference to the server ORB instance if the JNDI client does not provide an ORB instance. Typically, JNDI clients running in server processes use the server ORB instance; that is, they do not pass an ORB instance to the InitialContext constructor. The name server which is running in the server process sets a provider URL as a `java.lang.System` property to serve as the default provider URL for all JNDI clients in the process. This default provider URL is `corbaloc:rir:/NameServiceServerRoot`. This URL resolves to the server root context for that server. (The URL is equivalent to invoking `resolve_initial_references` on the ORB with a key of `NameServiceServerRoot`. The name server registers the server root context as an initial reference under that key.)

- **Understanding the legacy ORB protocol**

Previous versions of WebSphere Application Server used a different ORB implementation, which used a legacy protocol in contrast with the Interoperable Name Service (INS) protocol now used. This change has affected the implementation of the WebSphere Application Server initial context factory. **Certain types of pure clients can experience different behavior when getting initial JNDI contexts as compared to previous releases of WebSphere Application Server.** This behavior is discussed in more detail below.

The following ORB properties are used with the legacy ORB protocol for ORB initialization and are now deprecated:

- com.ibm.CORBA.BootstrapHost
- com.ibm.CORBA.BootstrapPort

The new INS ORB is different in a major respect, in that it exhibits no default behavior if no initial references are defined. In the legacy ORB, the bootstrap host and port values defaulted to localhost and 900. All initial references were obtained from the server running on the bootstrap host and port. So, if the ORB user provided no bootstrap host and port, all initial references are resolved from the server running on the local host at port 900. The INS ORB has no concept of bootstrap host or bootstrap port. All initial references are defined independently. That is, different initial references could resolve to different servers. If ORB.resolve\_initial\_references is invoked with a key such that the ORB is not initialized with an initial reference having that key, the call fails.

In previous releases of WebSphere Application Server, the initial context factory invoked resolve\_initial\_references on the ORB in the absence of any provider URL. This action succeeded if a name server at the default bootstrap host and port was running. Today, with the INS ORB, this would fail. (Actually, the ORB would fall back to the legacy protocol during the deprecation period, but when the legacy protocol is no longer supported, the operation would fail.) The initial context factory now uses a default provider URL of corbaloc:iiop:localhost:2809, and invokes string\_to\_object with the provider URL. This operation preserves the behavior that pure clients in previous releases experienced when they set no ORB bootstrap properties or provider URL. **However, this different initial context factory implementation changes the behavior experienced by certain legacy pure clients, which do not specify a provider URL:**

- Clients which set the ORB bootstrap properties listed above when getting an initial context.
- Clients which supply their own ORB instance to the InitialContext constructor.

There are two ways to circumvent this change of behavior:

- Always specify an IIOP type provider URL. This approach does not depend on the bootstrap host and port properties and continues to work when support for the bootstrap host and port properties is removed. For example, you can express bootstrap host and port property values of myHost and 2809, respectively, as corbaloc:iiop:myHost:2809.
- Use an rir type provider URL:
  - Specify corbaloc:rir:/NameServiceServerRoot if the ORB is initialized to use a WebSphere Application Server 5 server as the bootstrap server.
  - Specify corbaname:rir:/NameService#domain/legacyRoot if the ORB is initialized to use a WebSphere Application Server 4.0.x server as the bootstrap server.
  - Specify corbaloc:rir:/NameService if the ORB is initialized to use a server other than a WebSphere Application Server 5 or 4.0.x server as the bootstrap server.

URLs of this type are equivalent to invoking resolve\_initial\_references on the ORB with the specified key. If the bootstrap host and port properties are being used to initialize the ORB, this approach will not work when the bootstrap and host properties are no longer supported.

- **The InitialContext constructor search order for JNDI properties**

If the code snippet shown at the beginning of this section is executed by an application, the bootstrap server depends on the value of the property, java.naming.provider.url. If the property is not set (in server processes the default value is set as a system property), the default host of localhost and default port of 2809 are used as the address of the server from which to obtain the initial context. The JNDI specification describes where the InitialContext constructor looks for java.naming.provider.url property settings, but briefly, the property is picked up from the following places in the order shown:

1. The InitialContext constructor. This does not apply to the above example since the example uses the empty InitialContext constructor.

2. System environment. You can add JNDI properties to the system environment as an option on the Java command invocation and by program code. The recommended way to set the provider URL in the system environment is as an option supplied to the Java command invocation. Setting the provider URL in this manner is not temporal, so that getting a default initial context will always yield the same result. It is generally recommended that program code not set the provider URL property in the system environment because as a side-effect, this could adversely affect other, possibly unrelated, code running elsewhere in the same process.
3. `jndi.properties` file. There may be many `jndi.properties` files that are within the scope of the class loader in effect. All `jndi.properties` files are used for setting JNDI properties, but the provider URL setting is determined by the first `jndi.properties` file returned by the class loader.

## Example: Getting an initial context by setting the provider URL property

In general, JNDI clients should assume the correct environment is already configured so there is no need to explicitly set property values and pass them to the `InitialContext` constructor. However, a JNDI client may need to access a name space other than the one identified in its environment. In this case, it is necessary to explicitly set the `java.naming.provider.url` (provider URL) property used by the `InitialContext` constructor. A provider URL contains bootstrap server information that the initial context factory can use to obtain an initial context. Any property values passed in directly to the `InitialContext` constructor take precedence over settings of those same properties found elsewhere in the environment.

You can use two different provider URL forms with WebSphere Application Server's initial context factory:

- A CORBA object URL (new for J2EE 1.3)
- An IIOP URL

CORBA object URLs are more flexible than IIOP URLs and are the recommended URL format to use. CORBA object URLs are part of the OMG CosNaming Interoperable Naming Specification. A corbaname URL, for example, can include initial context and lookup name information and can be used as a lookup name without the need to explicitly obtain another initial context. The IIOP URLs are the legacy JNDI format, but are still supported by the WebSphere Application Server initial context factory.

The following examples illustrate the use of these URLs.

### Using a CORBA object URL

This example shows a CORBA object URL.

```
...
import java.util.Hashtable;
import javax.naming.Context;
import javax.naming.InitialContext;
...
Hashtable env = new Hashtable();
env.put(Context.INITIAL_CONTEXT_FACTORY,
 "com.ibm.websphere.naming.WsnInitialContextFactory");
env.put(Context.PROVIDER_URL, "corbaloc:iiop:myhost.mycompany.com:2809");
Context initialContext = new InitialContext(env);
...
```

### Using a CORBA object URL with multiple name server addresses

CORBA object URLs can contain more than one bootstrap address. You can use this feature when attempting to obtain an initial context from a server cluster. You can specify the bootstrap addresses for all servers in the cluster in the URL. The operation succeeds if at least one of the servers is running, eliminating a single point of failure. There is no guarantee of any particular order in which the address list will be processed. For example, the second bootstrap address may be used to obtain the initial context even though the server at the first bootstrap address in the list is available.

Multiple-address provider URLs should only contain the bootstrap addresses of members of the same cluster. Otherwise, incorrect behavior may occur.

An example of a corbaloc URL with multiple addresses follows.

```
...
import java.util.Hashtable;
import javax.naming.Context;
import javax.naming.InitialContext;
...
Hashtable env = new Hashtable();
env.put(Context.INITIAL_CONTEXT_FACTORY,
 "com.ibm.websphere.naming.WsnInitialContextFactory");
// All of the servers in the provider URL below are members of
// the same cluster.
env.put(Context.PROVIDER_URL,
 "corbaloc:myhost1:9810,myhost1:9811,myhost2:9810");
Context initialContext = new InitialContext(env);
...
```

## Using a CORBA object URL from a non-WebSphere Application Server JNDI implementation

Initial context factories for CosNaming JNDI plug-in implementations other than the WebSphere Application Server initial context factory most likely obtain an initial context using the object key, `NameService`. When you use such a context factory to obtain an initial context from a WebSphere Application Server name server, the initial context is the cell root context. Since system artifacts such as EJB homes associated with a server are bound under the server's server root context, names used in JNDI operations must be qualified. If you want to use relative names, ensure your initial context is the server root context under which the target object is bound. In order to make the server root context the initial context, specify a corbaloc provider URL with an object key of `NameServiceServerRoot`.

This example shows a CORBA object type URL from a non-WebSphere Application Server JNDI implementation. This example assumes full CORBA object URL support by the non-WebSphere Application Server JNDI implementation. The object key of `NameServiceServerRoot` is specified so that the initial context will be the specified server's server root context.

```
...
import java.util.Hashtable;
import javax.naming.Context;
import javax.naming.InitialContext;
...
Hashtable env = new Hashtable();
env.put(Context.INITIAL_CONTEXT_FACTORY,
 "com.somecompany.naming.TheirInitialContextFactory");
env.put(Context.PROVIDER_URL,
 "corbaname:iiop:myhost.mycompany.com:9810/NameServiceServerRoot");
Context initialContext = new InitialContext(env);
...
```

If qualified names are used, you can use the default key of `NameService`.

## Using an IIOP URL

The IIOP type of URL is a legacy format which is not as flexible as CORBA object URLs. However, URLs of this type are still supported. The following example shows an IIOP type URL as the provider URL.

```
...
import java.util.Hashtable;
import javax.naming.Context;
import javax.naming.InitialContext;
...
Hashtable env = new Hashtable();
env.put(Context.INITIAL_CONTEXT_FACTORY,
 "com.ibm.websphere.naming.WsnInitialContextFactory");
env.put(Context.PROVIDER_URL, "iiop://myhost.mycompany.com:2809");
Context initialContext = new InitialContext(env);
...
```

## Example: Setting the provider URL property to select a different root context as the initial context

Each server contains its own server root context, and, when bootstrapping to a server, the server root is the default initial JNDI context. Most of the time, this default is the desired initial context, since system artifacts such as EJB homes are bound there. However, other root contexts exist, which can contain bindings of interest. It is possible to specify a provider URL to select other root contexts.

### Selecting the initial root context with a CORBA object URL

There are several object keys registered with the bootstrap server that you can use to select the root context for the initial context. To select a particular root context with a CORBA object URL object key, set the object key to the corresponding value. The default object key is NameService. Using JNDI yields the server root context. A table that lists the different root contexts and their corresponding object key follows:

Root Context	CORBA Object URL Object Key
Server Root	NameServiceServerRoot
Cell Persistent Root	NameServiceCellPersistentRoot
Cell Root	NameServiceCellRoot
Node Root	NameServiceNodeRoot

The following example shows the use of a corbaloc URL with the object key set to select the cell persistent root context as the initial context.

```
...
import java.util.Hashtable;
import javax.naming.Context;
import javax.naming.InitialContext;
...
Hashtable env = new Hashtable();
env.put(Context.INITIAL_CONTEXT_FACTORY,
 "com.ibm.websphere.naming.WsnInitialContextFactory");
env.put(Context.PROVIDER_URL,
 "corbaloc:iiop:myhost.mycompany.com:2809/NameServiceCellPersistentRoot");
Context initialContext = new InitialContext(env);
...
```

### Selecting the initial root context with the name space root property

You can also select the initial root context by passing a name space root property setting to the InitialContext constructor. Generally, the object key setting described above is sufficient. Sometimes a property setting is preferable. For example, you can set the root context property on the Java invocation to make which server root is being used as the initial context transparent to the application. The default server root property setting is defaultroot, which yields the server root context.

Root Context	Name Space Root Property Value
Server Root	bootstrapserverroot
Cell Persistent Root	cellpersistentroot
Cell Root	cellroot
Node Root	bootstrapnoderoot

The initial context factory ignores the name space root property if the provider URL contains an object key other than NameService.

The following example shows use of the name space root property to select the cell persistent root context as the initial context. Note that available constants are used instead of hard-coding the property name and value.



```

...
import java.util.Hashtable;
import javax.naming.Context;
import javax.naming.InitialContext;
import com.ibm.websphere.naming.PROPS;
...
Hashtable env = new Hashtable();
env.put(Context.INITIAL_CONTEXT_FACTORY,
 "com.ibm.websphere.naming.WsnInitialContextFactory");
env.put(Context.PROVIDER_URL, "corbaloc:iiop:myhost.mycompany.com:2809");
env.put(PROPS.NAME_SPACE_ROOT, PROPS.NAME_SPACE_ROOT_CELL_PERSISTENT);
Context initialContext = new InitialContext(env);
...

```

## Example: Looking up an EJB home with JNDI

Most applications which use JNDI run in a container. Some do not. The name used to look up an object depends on whether or not the application is running in a container. The examples below show lookups from each type of application. Sometimes it is more convenient for an application to use a corbaname URL as the lookup name. Container-based JNDI clients and thin Java clients can use a corbaname URL. An example of a lookup with a corbaname URL is also included in this section.

### JNDI lookup from an application running in a container

Applications that run in a container can use `java:` lookup names. Lookup names of this form provide a level of indirection such that the lookup name used to look up an object is not dependent on the object's name as it is bound in the name server's name space. The deployment descriptors for the application provide the mapping from the `java:` name and the name server lookup name. The container sets up the `java:` name space based on the deployment descriptor information so that the `java:` name is correctly mapped to the corresponding object.

The following example shows a lookup of an EJB home. The actual home lookup name is determined by the application's deployment descriptors.

```

// Get the initial context as shown in a previous example
...
// Look up the home interface using the JNDI name
try {
 java.lang.Object ejbHome =
 initialContext.lookup(
 "java:comp/env/com/mycompany/accounting/AccountEJB");
 accountHome = (AccountHome)javax.rmi.PortableRemoteObject.narrow(
 (org.omg.CORBA.Object) ejbHome, AccountHome.class);
}
catch (NamingException e) { // Error getting the home interface
 ...
}

```

### JNDI lookup from an application that does not run in a container

Applications that do not run in a container cannot use `java:` lookup names because it is the container which sets the `java:` name space up for the application. Instead, an application of this type must look the object up directly from the name server. Each application server contains a name server. System artifacts such as EJB homes are bound relative to the server root context in that name server. The various name servers are federated by means of a system name space structure. The recommended way to look up objects on different servers is to qualify the name so that the name resolves from any initial context in the cell. If a relative name is used, the initial context must be the same server root context as the one under which the object is bound. The form of the qualified name depends on whether the qualified name is a topology-based name or a fixed name. A topology based name depends on whether the object resides in a single server or a server cluster. Examples of each form of qualified name follow.

- **Topology-based qualified names**



Topology-based qualified names traverse through the system name space to the server root context context under which the target object is bound. A topology-based qualified name resolves from any initial context in the cell. The topology-based qualified name depends on whether the object resides on a single server or server cluster. Examples of each lookup follow.

### Single server

The following example shows a lookup of an EJB home that is running in the single server, MyServer, configured in the node, Node1.

```
// Get the initial context as shown in a previous example
// Using the form of lookup name below, it doesn't matter which
// server in the cell is used to obtain the initial context.
...
// Look up the home interface using the JNDI name
try {
 java.lang.Object ejbHome = initialContext.lookup(
 "cell/nodes/Node1/servers/MyServer/com/mycompany/accounting/AccountEJB");
 accountHome = (AccountHome)javadoc.rmi.PortableRemoteObject.narrow(
 (org.omg.CORBA.Object) ejbHome, AccountHome.class);
}
catch (NamingException e) { // Error getting the home interface
 ...
}
```

### Server cluster

The example below shows a lookup of an EJB home which is running in the cluster, MyCluster. The name can be resolved if any of the cluster members is running.

```
// Get the initial context as shown in a previous example
// Using the form of lookup name below, it doesn't matter which
// server in the cell is used to obtain the initial context.
...
// Look up the home interface using the JNDI name
try {
 java.lang.Object ejbHome = initialContext.lookup(
 "cell/clusters/MyCluster/com/mycompany/accounting/AccountEJB");
 accountHome = (AccountHome)javadoc.rmi.PortableRemoteObject.narrow(
 (org.omg.CORBA.Object) ejbHome, AccountHome.class);
}
catch (NamingException e) { // Error getting the home interface
 ...
}
```

- **Fixed qualified names**

If the target object has a cell-scoped fixed name defined for it, you can use its qualified form instead of the topology-based qualified name. Even though the topology-based name works, the fixed name does not change with the specific cell topology or with the movement of the target object to a different server. An example lookup with a qualified fixed name is shown below.

```
// Get the initial context as shown in a previous example
// Using the form of lookup name below, it doesn't matter which
// server in the cell is used to obtain the initial context.
...
// Look up the home interface using the JNDI name
try {
 java.lang.Object ejbHome = initialContext.lookup(
 "cell/persistent/com/mycompany/accounting/AccountEJB");
 accountHome = (AccountHome)javadoc.rmi.PortableRemoteObject.narrow(
 (org.omg.CORBA.Object) ejbHome, AccountHome.class);
}
catch (NamingException e) { // Error getting the home interface
 ...
}
```

## JNDI lookup with a corbaname URL

A corbaname can be useful at times as a lookup name. If, for example, the target object is not a member of the federated name space and cannot be located with a qualified name, a corbaname can be a convenient way to look up the object. A lookup with a corbaname URL follows.

```
// Get the initial context as shown in a previous example
...
// Look up the home interface using a corbaname URL
try {
 java.lang.Object ejbHome = initialContext.lookup(
 "corbaname:iiop:someHost:2809#com/mycompany/accounting/AccountEJB");
 accountHome = (AccountHome)javax.rmi.PortableRemoteObject.narrow(
 (org.omg.CORBA.Object) ejbHome, AccountHome.class);
}
catch (NamingException e) { // Error getting the home interface
 ...
}
```

## Example: Looking up a JavaMail session with JNDI

The example below shows a lookup of a JavaMail resource. The actual lookup name is determined by the application's deployment descriptors.

```
// Get the initial context as shown above
...
Session session =
 (Session) initialContext.lookup("java:comp/env/mail/MailSession");
```

## JNDI interoperability considerations

This section explains considerations to take into account when interoperating with previous releases of WebSphere Application Server and with non-WebSphere Application Server JNDI clients. Also, the way resources from MQSeries must be bound to the name space has changed and is described below.

### Interoperability with previous WebSphere Application Server Releases

- **EJB clients running on WebSphere Application Server V3.5 or V4.0 accessing EJB applications running on WebSphere Application Server V5**

Applications migrated from previous versions of WebSphere Application Server may still have clients still running in a previous release. The default initial JNDI context for EJB clients running on previous versions of WebSphere Application Server is the cell persistent root (legacy root). The home for an enterprise bean deployed in version 5 is bound to its server's server root context. In order for the EJB lookup name for down-level clients to remain unchanged, configure a binding for the EJB home under the cell persistent root.

**Note:** EJB clients running in version 3.5 must be running in version 3.5.5 or above, or in version 3.5.3 or 3.5.4 with e-fix PQ51387 installed.

- **EJB clients running on WebSphere Application Server V5 accessing EJB applications running on WebSphere Application Server V3.5 or V4.0 servers**

The default initial context for a WebSphere Application Server v3.5 or v4.0 server is the correct initial context. Simply look up the JNDI name under which the EJB home is bound.

**Note:** To enable WebSphere Application Server V5 clients to access version 3.5.x and 4.0.x servers, the down-level installations must have e-fix PQ60074 installed.

### EJB clients running in an environment other than WebSphere Application Server accessing EJB applications running on WebSphere Application Server V5 servers

When an EJB application running in WebSphere Application Server V5 is accessed by a non-WebSphere Application Server EJB client, the JNDI initial context factory is presumed to be a non-WebSphere

Application Server implementation. In this case, the default initial context will be the cell root. If the JNDI service provider being used supports CORBA object URLs, the corbaname format can be used to look up the EJB home. The construction of the stringified name depends on whether the object is installed on a single server or cluster, as shown below.

- **Single server**

```
initialContext.lookup(
 "corbaname:iiop:myHost:2809#cell/nodes/node1/servers/server1/myEJB");
```

According to the URL above, the bootstrap host and port are myHost and 2809, and the enterprise bean is installed in a server **server1** in node **node1** and bound in that server under the name **myEJB**.

- **Server cluster**

```
initialContext.lookup(
 "corbaname:iiop:myHost:2809#cell/clusters/myCluster/myEJB");
```

According to the URL above, the bootstrap host and port are **myHost** and **2809**, and the enterprise bean is installed in a server cluster named **myCluster** and bound in that cluster under the name **myEJB**.

The above lookup will work with any name server bootstrap host and port configured in the same cell.

The above lookup will also work if the bootstrap host and port belongs to a member of the cluster itself.

To avoid a single point of failure, the bootstrap server host and port for each cluster member could be listed in the URL as follows:

```
initialContext.lookup(
 "corbaname:iiop:host1:9810,host2:9810#cell/clusters/myCluster/myEJB");
```

The name prefix **cell/clusters/myCluster/** is not necessary if bootstrapping to the cluster itself, but it will work. The prefix is needed, however, when looking up enterprise beans in other clusters. Name bindings under the **clusters** context are implemented on the name server to resolve to the server root of a running cluster member during a lookup; thus avoiding a single point of failure.

- **Without CORBA object URL support**

If the JNDI initial context factory being used does not support CORBA object URLs, the initial context can be obtained from the server, and the lookup can be performed on the initial context as follows:

```
Hashtable env = new Hashtable();
env.put(CONTEXT.PROVIDER_URL, "iiop://myHost:2809");
Context ic = new InitialContext(env);
Object o = ic.lookup("cell/clusters/myCluster/myEJB");
```

## Binding resources from MQSeries 5.2

In previous releases of WebSphere Application Server, the MQSeries jmsadmin tool could be used bind resources to the name space. When used with a WebSphere Application Server V5 name space, the resource will be bound within a transient partition in the name space and will not persist past the life of the server process. Instead of binding the MQSeries resources with the jmsadmin tool, bind them from the WebSphere Application Server administrative console, under Resources in the left panel on the console

## JNDI caching

To increase the performance of JNDI operations, the WebSphere Application Server JNDI implementation employs caching to reduce the number of remote calls to the name server for lookup operations. For most cases, use the default cache setting.

When an InitialContext object is instantiated, an association is established between the InitialContext instance and a cache. The initial context and any contexts returned directly or indirectly from a lookup on the initial context are all associated with that same cache instance. By default, the association is based on the provider URL, in particular, the host name and port. The caller can specify the cache name to override this default behavior. A cache instance of a given name is shared by all instances of InitialContext configured to use a cache of that name which were created with the same context class loader in effect. Two EJB applications running in the same server will use their own cache instances, if they are using different context class loaders, even if the cache names are the same.

After an association between an InitialContext instance and cache is established, the association does not change. A `javax.naming.Context` object returned from a lookup operation inherits the cache association of the Context object on which the lookup was performed. Changing cache property values with the `Context.addToEnvironment()` or `Context.removeFromEnvironment()` method does not affect cache behavior. You can change properties affecting a given cache instance with each InitialContext instantiation.

A cache is restricted to a process and does not persist past the life of that process. A cached object is returned from lookup operations until either the max cache life for the cache is reached, or the max entry life for the object's cache entry is reached.

After this time, a lookup on the object causes the cache entry for the object to be refreshed. If a bind or rebind operation is executed on an object, the change is not reflected in any caches other than the one associated with the context from which the bind or rebind was issued. This scenario is most likely to happen when multiple processes are involved, since different processes do not share the same cache, and context objects in all threads in a process typically share the same cache instance for a given name service provider.

Usually, cached objects are relatively static entities, and objects becoming stale are not a problem. However, you can set timeout values on cache entries or on a cache so that cache contents are periodically refreshed.

## JNDI cache settings

Various cache property settings follow. Ensure that all property values are string values.

### **com.ibm.websphere.naming.jndicache.cachename**

The name of the cache to associate with an initial context instance can be specified with this property.

It is possible to create multiple InitialContext instances, each operating on the name space of a different name server. By default, objects from each bootstrap address are cached separately, since they each involve independent name spaces and name collisions could occur if they used the same cache. The provider URL specified when the initial context is created by default serves as the basis for the cache name. With this property, a JNDI client can specify a cache name. Valid options for cache names follow:

Valid options	Resulting cache behavior
<b>providerURL (default)</b>	Use the value for <code>java.naming.provider.url</code> property as the basis for the cache name. Cache names are based on the bootstrap host and port specified in the URL. The bootstrap host is normalized to a fully qualified name, if possible. For example, <code>"corbaname:iiop:server1:2809#some/starting/context"</code> and <code>"corbaloc:iiop://server1"</code> are normalized to the same cache name. If no provider URL is specified, a default cache name is used.
<b>Any string</b>	Use the specified string as the cache name. You can use any arbitrary string with a value other than <code>"providerURL"</code> as a cache name.

### **com.ibm.websphere.naming.jndicache.cacheobject**

Turn caching on or off and clear an existing cache with this property.

By default, when an InitialContext is instantiated, it is associated with an existing cache or, if one does not exist, a new one is created. An existing cache is used with its existing contents. In some circumstances, this behavior is not desirable. For example, when objects that are looked up change frequently, they can become stale in the cache. Other options are available. The following table lists these other options along with the corresponding property value.

Valid values	Resulting cache behavior
--------------	--------------------------

<b>populated (default)</b>	Use a cache with the specified name. If the cache already exists, leave existing cache entries in the cache; otherwise, create a new cache.
<b>cleared</b>	Use a cache with the specified name. If the cache already exists, clear all cache entries from the cache; otherwise, create a new cache.
<b>none</b>	Do not cache. If this option is specified, the cache name is irrelevant. Therefore, this option will not disable a cache that is already associated with other InitialContext instances. The InitialContext that is instantiated is not associated with any cache.

### **com.ibm.websphere.naming.jndicache.maxcachelife**

Impose a limit to the age of a cache with this property.

By default, cached objects remain in the cache for the life of the process or until cleared with the `com.ibm.websphere.naming.jndicache.cacheobject` property set to "cleared". This property enables a JNDI client to set the maximum life of a cache. This property differs from the `maxentrylife` property (below) in that the entire cache is cleared when the cache lifetime is reached. The table below lists the various `maxcachelife` values and their affect on cache behavior:

<b>Valid options</b>	<b>Resulting cache behavior</b>
<b>0 (default)</b>	Make the cache lifetime unlimited.
<b>Positive integer</b>	Set the maximum lifetime of the entire cache, in minutes, to the specified value. When the maximum lifetime for the cache is reached, the next attempt to read any entry from the cache causes the cache to be cleared

### **com.ibm.websphere.naming.jndicache.maxentrylife**

Impose a limit to the age of individual cache entries with this property.

By default, cached objects remain in the cache for the life of the process or until cleared with the `com.ibm.websphere.naming.jndicache.cacheobject` property set to `cleared`. This property enables a JNDI client to set the maximum lifetime of individual cache entries. This property differs from the `maxcachelife` property in that individual entries are refreshed individually as their maximum lifetime reached. This might avoid any noticeable change in performance that might occur if the whole cache is cleared at once. The table below lists the various `maxentrylife` values and their effect on cache behavior:

<b>Valid options</b>	<b>Resulting cache behavior</b>
<b>0 (default)</b>	Lifetime of cache entries is unlimited.
<b>Positive integer</b>	Set the maximum lifetime of individual cache entries, in minutes, to the specified value. When the maximum lifetime for an entry is reached, the next attempt to read the entry from the cache causes the individual cache entry to refresh.

## **Example: Controlling JNDI cache behavior from a program**

Following are examples that illustrate how you can use JNDI cache properties to achieve the desired cache behavior. Cache properties take effect when an InitialContext object is constructed.

```
import java.util.Hashtable;
import javax.naming.InitialContext;
import javax.naming.Context;
import com.ibm.websphere.naming.PROPS;
```

```
/*****
```

```
Caching discussed in this section pertains to the WebSphere Application
Server initial context factory. Assume the property,
java.naming.factory.initial, is set to
"com.ibm.websphere.naming.WsnInitialContextFactory" as a
java.lang.System property.
```

```

*****/

Hashtable env;
Context ctx;

// To clear a cache:

env = new Hashtable();
env.put(PROPS.JNDI_CACHE_OBJECT, PROPS.JNDI_CACHE_OBJECT_CLEARED);
ctx = new InitialContext(env);

// To set a cache's maximum cache lifetime to 60 minutes:

env = new Hashtable();
env.put(PROPS.JNDI_CACHE_MAX_LIFE, "60");
ctx = new InitialContext(env);

// To turn caching off:

env = new Hashtable();
env.put(PROPS.JNDI_CACHE_OBJECT, PROPS.JNDI_CACHE_OBJECT_NONE);
ctx = new InitialContext(env);

// To use caching and no caching:

env = new Hashtable();
env.put(PROPS.JNDI_CACHE_OBJECT, PROPS.JNDI_CACHE_OBJECT_POPULATED);
ctx = new InitialContext(env);
env.put(PROPS.JNDI_CACHE_OBJECT, PROPS.JNDI_CACHE_OBJECT_NONE);
Context noCacheCtx = new InitialContext(env);

Object o;

// Use caching to look up home, since the home should rarely change.
o = ctx.lookup("com/mycom/MyEJBHome");
// Narrow, etc. ...

// Do not use cache if data is volatile.
o = noCacheCtx.lookup("com/mycom/VolatileObject");
// ...

```

## JNDI name syntax

JNDI name syntax is the default syntax and is suitable for typical JNDI clients.

This syntax includes the following special characters: forward slash (/) and backslash (\). Components in a name are delimited by a forward slash. The backslash is used as the escape character. A forward slash is interpreted literally if it is escaped, that is, preceded by a backslash. Similarly, a backslash is interpreted literally if it is escaped.

## INS name syntax

INS syntax is designed for JNDI clients that need to interoperate with CORBA applications.

The INS syntax allows a JNDI client to make the proper mapping to and from a CORBA name. INS syntax is very similar to the JNDI syntax with the additional special character, dot (.). Dots are used to delimit the `id` and `kind` fields in a name component. A dot is interpreted literally when it is escaped. Only one unescaped dot is allowed in a name component. A name component with a non-empty `id` field and empty `kind` field is represented with only the `id` field value and must not end with an unescaped dot. An empty name component (empty `id` and empty `kind` field) is represented with a single unescaped dot. An empty string is not a valid name component representation.

## JNDI to CORBA name mapping considerations

WebSphere Application Server name servers are an implementation of the CORBA CosNaming interface. WebSphere Application Server provides a JNDI implementation which you can use to access CosNaming name servers through the JNDI interface. Issues can exist when mapping JNDI name strings to and from CORBA names.

Each component in a CORBA name consists of an `id` and `kind` field, but a JNDI name component consists of no such fields. Each component in a JNDI name is atomic. Typical JNDI clients do not need to make a distinction between the `id` and `kind` fields of a name component, or know how JNDI name strings map to CORBA names. JNDI clients of this sort can use the JNDI syntax described below. When a name is parsed according to JNDI syntax, each name component is mapped to the `id` field of the corresponding CORBA name component. The `kind` field always has an empty value. This basic syntax is the least obtrusive to the JNDI client in that it has the fewest special characters. However, you cannot represent with this syntax a CORBA name with a non-empty `kind` field. This restriction can prevent EJB applications from interoperating with CORBA applications.

Some clients, however must interoperate with CORBA applications which use CORBA names with non-empty `kind` fields. These JNDI clients must make a distinction between `id` and `kind` so that JNDI names are correctly mapped to CORBA names, particularly when the CORBA names contain components with non-null `kind` fields. Such JNDI clients can use the INS name syntax. With its additional special character, you can use INS to represent any CORBA name. Use of this syntax is not recommended unless it is necessary, because this syntax is more restrictive from the JNDI client's perspective in that the JNDI client must be aware that name components with multiple unescaped dots are syntactically invalid. INS name syntax is part of the OMG CosNaming Interoperable Naming Specification.

### Example: Setting the syntax used to parse name strings

JNDI clients which must interoperate with CORBA applications may need to use INS name syntax to represent names in string format. The name syntax property may be passed to the `InitialContext` constructor through its parameter, in the System properties, or in a `jndi.properties` file. The initial context and any contexts looked up from that initial context will parse name strings based on the specified syntax.

The following example shows how to set the name syntax to make the initial context parse name strings according to INS syntax.

```
...
import java.util.Hashtable;
import javax.naming.Context;
import javax.naming.InitialContext;
import com.ibm.websphere.naming.PROPS; // WebSphere naming constants
...
Hashtable env = new Hashtable();
env.put(Context.INITIAL_CONTEXT_FACTORY,
 "com.ibm.websphere.naming.WsnInitialContextFactory");
env.put(Context.PROVIDER_URL, ...);
env.put(PROPS.NAME_SYNTAX, PROPS.NAME_SYNTAX_INS);
Context initialContext = new InitialContext(env);
// The following name maps to a CORBA name component as follows:
// id = "a.name", kind = "in.INS.format"
// The unescaped dot is used as the delimiter.
// Escaped dots are interpreted literally.
java.lang.Object o = initialContext.lookup("a\\.name.in\\.INS\\.format");
...
```



---

## Developing applications that use CosNaming (CORBA Naming interface)

CORBA clients can perform naming operations on WebSphere name servers through the CosNaming interface. The following examples show how to obtain an ORB instance and an initial context as well as how to look up an EJB home.

**Note:** To enable WebSphere Application Server V5 clients to access Versions 3.5.x and 4.0.x servers, the earlier installations must have e-fix PQ60074 installed.

1. Get an initial context
2. Perform desired CosNaming operations

### Example: Getting an initial context with CosNaming

In the WebSphere Application Server, an initial context is obtained from a bootstrap server. The address for the bootstrap server consists of a host and port. To get an initial context, you must know the host and port for the server that is used as the bootstrap server.

Obtaining an initial context consists of two basic steps:

1. Obtain an ORB reference
2. Invoke a method on the ORB to obtain the initial reference

These steps are now explained in more detail.

### Obtaining an ORB reference

Pure CosNaming clients, that is clients that are not running in a server process, must create and initialize an ORB instance with which to obtain the initial context. CosNaming clients which run in server processes can obtain a reference to the server ORB with a JNDI lookup. The following examples illustrate how to create and initialize a client ORB and how to obtain a server ORB reference.

### Creating a client ORB instance

To create an ORB instance, invoke the static method, `org.omg.CORBA.ORB.init`. The `init` method requires a property set to the name of the ORB class you want to instantiate. An ORB implementation with the class name `com.ibm.CORBA.iop.ORB` is included with the WebSphere Application Server. The WebSphere Application Server ORB recognizes additional properties with which you can specify initial references.

The basic steps for creating an ORB are as follows:

1. Create a Properties object.
2. Set the ORB class property to WebSphere Application Server's ORB class.
3. If the bootstrap server is INS-compliant, set the initial reference properties. If the bootstrap server is not INS-compliant (meaning, WebSphere Application Server v4.0.x or earlier), set bootstrap host and port for bootstrap server.
4. Invoke `ORB.init`, passing in the Properties object.

### Usage scenario

```
...
import java.util.Properties;
import org.omg.CORBA.ORB;
...
Properties props = new Properties();
props.put("org.omg.CORBA.ORBClass", "com.ibm.CORBA.iop.ORB");
props.put("com.ibm.CORBA.ORBInitRef.NameService",
 "corbaloc:iiop:myhost.mycompany.com:2809/NameService");
```

```

props.put("com.ibm.CORBA.ORBInitRef.NameServiceServerRoot",
"corbaloc:iiop:myhost.mycompany.com:2809/NameServiceServerRoot");
ORB _orb = ORB.init((String[])null, props);
...

```

Notice the initial reference definitions for NameService and NameServiceServerRoot. The initial context returned for NameService depends on the type of bootstrap server. The key NameServiceServerRoot is a key introduced in WebSphere Application Server V5. For more information on initial contexts, see the section Initial Contexts.

**Note:** The properties com.ibm.CORBA.BootstrapHost and com.ibm.CORBA.BootstrapPort are deprecated. They are needed, however, to connect to WebSphere Application Servers of Version 4.0.x or earlier. The default bootstrap host is the local host and the default port is 2809.

### Obtaining a reference to the server ORB

CosNaming clients which run in a server process can obtain a reference to the server ORB with a JNDI lookup on a java: name, shown as follows:

#### Usage scenario

```

...
import javax.naming.Context;
import javax.naming.InitialContext;
import org.omg.CORBA.ORB;
...
Context initialContext = new InitialContext();
ORB orb = (ORB) initialContext.lookup("java:comp/ORB");
...

```

### Using an ORB reference to get an initial naming reference

There are two basic ways to get an initial CosNaming context. Both ways involve an ORB method invocation. The first way is to invoke the resolve\_initial\_references method on the ORB with an initial reference key. For this call to work, the ORB must be initialized with an initial reference for that key. The other way is to invoke the string\_to\_object method on the ORB, passing in a CORBA object URL with the host and port of the bootstrap server. The following examples illustrate both approaches.

#### Invoking resolve\_initial\_references

Once an ORB reference is obtained, invoke the resolve\_initial\_references method on the ORB to obtain a reference to the initial context. The following code example invokes resolve\_initial\_reference on an ORB reference.

#### Usage scenario

```

...
import org.omg.CORBA.ORB;
import org.omg.CosNaming.NamingContextExt;
import org.omg.CosNaming.NamingContextExtHelper;
...
// Obtain ORB reference as shown in examples earlier in this section
...
org.omg.CORBA.Object obj = _orb.resolve_initial_references("NameService");
NamingContextExt initCtx = NamingContextExtHelper.narrow(obj);
...

```

Note that the key NameService is passed to the resolve\_initial\_references method. Other initial context keys are registered in WebSphere Application Servers. For example, NameServiceServerRoot can be used to obtain a reference to the server root context in the bootstrap name server. For more information on the initial contexts registered in server ORBs, please see the section Initial Contexts.

## Invoking string\_to\_object with a CORBA object URL

You can use an INS-compliant ORB to obtain an initial context even if the ORB is not initialized with any initial references or bootstrap properties, or if those property settings are for a different server than the name server from which you want to obtain the initial context. To obtain an initial context by explicitly specifying the bootstrap name server, invoke the `string_to_object` method on the ORB, passing in a CORBA object URL which contains the bootstrap server host and port.

The code in the example below invokes the `string_to_object` method on an existing ORB reference, passing in a CORBA object URL which identifies the desired initial context.

### Usage scenario

```
...
import org.omg.CORBA.ORB;
import org.omg.CosNaming.NamingContextExt;
import org.omg.CosNaming.NamingContextExtHelper;
...
// Obtain ORB reference as shown in examples earlier in this section
...
org.omg.CORBA.Object obj =
 orb.string_to_object("corbaloc:iiop:myhost.mycompany.com:2809/NameService");
NamingContextExt initCtx = NamingContextExtHelper.narrow(obj);
...
```

Note that the key `NameService` is used in the `corbaloc` URL. Other initial context keys are registered in WebSphere Application Servers. For example, you can use `NameServiceServerRoot` to obtain a reference to the server root context in the bootstrap name server.

## Using an existing ORB and invoking string\_to\_object with a CORBA object URL with multiple name server addresses to get an initial context

CORBA object URLs can contain more than one bootstrap server address. Use this feature when attempting to obtain an initial context from a server cluster. You can specify the bootstrap server addresses for all servers in the cluster in the URL. The operation will succeed if at least one of the servers is running, eliminating a single point of failure. There is no guarantee of any particular order in which the address list will be processed. For example, the second bootstrap server address may be used to obtain the initial context even though the first bootstrap server in the list is available. An example of a `corbaloc` URL with multiple addresses follows.

```
...
import org.omg.CORBA.ORB;
import org.omg.CosNaming.NamingContextExt;
import org.omg.CosNaming.NamingContextExtHelper;
...
// Assume orb is an existing ORB instance
org.omg.CORBA.Object obj = orb.string_to_object(
 "corbaloc::myhost1:9810,:myhost1:9811,:myhost2:9810/NameService");
NamingContextExt initCtx = NamingContextExtHelper.narrow(obj);
...
```

## Example: Looking up an EJB home with CosNaming

You can look up an EJB home or other CORBA object from a WebSphere Application Server name server through the CORBA `CosNaming` interface. You can invoke `resolve` or `resolve_str` on the initial context, or you can invoke `string_to_object` on the ORB. You can use a qualified name so that the name resolves regardless of which name server the lookup is executed on, or use an unqualified name that only resolves from the server root context on the name server that actually contains the object binding. (The qualified name traverses the federated system name space to the specified server root context.)

## Qualified and unqualified names

Each application server contains a name server. System artifacts such as EJB homes are bound in that name server. The various name servers are federated by means of a system name space structure. The recommended way to look up objects on different servers is to use a qualified name. A qualified name can be a topology-based name, based on the name of the cluster or single server and node that contains the object. You can define fixed qualified names for objects. With qualified names, you can look up objects residing on different servers from the same initial context by traversing the system name space structure. Alternatively, you can use an unqualified name, but an unqualified name will only resolve using the name server associated with the object's application server.

## CosNaming.resolve (and resolve\_str) vs. ORB.string\_to\_object

If you have an initial context from any name server in a WebSphere Application Server cell, you can look up any CORBA object with a qualified name. You do not need additional host and port information for the target object's name server.

Alternatively, you can look up an object by invoking `string_to_object` on the ORB, passing in a corbaname URL. Typically, an IIO type URL is specified, so the bootstrap address information required for an initial context must be contained in the URL. You can use a qualified or unqualified stringified name, but an unqualified name resolves only if the initial context is from the name server in which the object is bound.

The following examples show CosNaming resolve operations using qualified topology-based lookup names and an unqualified lookup name.

## CosNaming resolve operation using a qualified name

The topology-based qualified name for an object depends on whether the object is bound in a single server or a server cluster. Examples of each follow.

### Single Server

The following example shows the lookup of an EJB home that is running in a single server. The enterprise bean that is being looked up is running in the server, `MyServer`, on the node, `Node1`.

```
// Get the initial context as shown in the previous example
// Using the form of lookup name below, it doesn't matter which
// server in the cell is used to obtain the initial context.
...
// Look up the home interface using the name under which the EJB home is bound
org.omg.CORBA.Object ejbHome = initialContext.resolve_str(
 "cell/nodes/Node1/servers/MyServer/mycompany/accounting/AccountEJB");
accountHome =
 (AccountHome)javax.rmi.PortableRemoteObject.narrow(ejbHome, AccountHome.class);
```

### Server Cluster

The following example shows a lookup of an EJB home that is running in a cluster. The enterprise bean being that is looked up is running in the cluster, `Cluster1`. The name can be resolved if any of the cluster members is running.

### Usage scenario

```
// Get the initial context as shown in the previous example
// Using the form of lookup name below, it doesn't matter which
// server in the cell is used to obtain the initial context.
...
// Look up the home interface using the name under which the EJB home is bound
org.omg.CORBA.Object ejbHome = initialContext.resolve_str(
```

```

 "cell/clusters/Cluster1/mycompany/accounting/AccountEJB");
accountHome =
 (AccountHome)javax.rmi.PortableRemoteObject.narrow(ejbHome, AccountHome.class);

```

## ORB string\_to\_object operation using an unqualified stringified name

If the resolve operation is being performed on the name server that contains the object, the system name space does not need to be traversed, and you can use an unqualified lookup name. Note that this name does not resolve on other name servers. If an unqualified name is provided, the object key must be `NameServiceServerRoot` so that the correct initial context is selected. If a qualified name is provided, you can use the default key of `NameService`.

The following example shows a lookup of an EJB home. The enterprise bean that is being looked up is bound on the name server running on the host `myHost` on port `2809`. Note the object key of `NameServiceServerRoot`.

```

// Assume orb is an existing ORB instance
...
// Look up the home interface using the name under which the EJB home is bound
org.omg.CORBA.Object ejbHome = orb.string_to_object(
 "corbaname:iiop:myHost:2809/NameServiceServerRoot#mycompany/accounting");
accountHome =
 (AccountHome)javax.rmi.PortableRemoteObject.narrow(ejbHome, AccountHome.class);

```

---

## Configured name bindings

Administrators can configure bindings into the name space. A configured binding is different from a programmatic binding in that the system creates the binding every time a server is started, even if the target context is in a transient partition.

Administrators can add name bindings to the name space through the configuration. Name servers add these configured bindings to the name space view, by reading the configuration data for the bindings. Configuring bindings is an alternative to creating the bindings from a program. Configured bindings have the advantage of being created each time a server starts, even when the binding is created in a transient partition of the name space. Cell-scoped configured bindings provide interoperability with JNDI clients running on previous versions of WebSphere Application Server. Additionally, you can configure cell-scoped bindings to create a fixed qualified name for server objects.

### Scope

You can configure a binding at one of the following three scopes: cell, node, or server. Cell-scoped bindings are created under the cell persistent root context. Node-scoped bindings are created under the node persistent root context for the specified node. Server-scoped bindings are created under the server root context for the selected server. If the target server of a server-scoped binding is a cluster, the binding is created under the server root context of each cluster member.

**Note:** The term *server* includes clusters and can be used interchangeably with the term *cluster* with respect to configured bindings. When applied to a cluster, a server-scoped binding is created in the server root for all member servers.

The scope you select for new bindings depends on how the binding is to be used. For example, if the binding is not specific to any particular node or server, or if you do not want the binding to be associated with any specific node or server, a cell-scoped binding is a suitable scope. Defining fixed names for enterprise beans to create fixed qualified names is just such an application. If a binding is to be used only by clients of an application running on a particular server, or if you want to configure a binding with the same name on different servers which resolve to different objects, a server-scoped binding would be appropriate. Note that two servers can have configured bindings with the same name but resolve to different objects. At the cell scope, only one binding with a given name can exist.

## Intermediate Contexts

Intermediate contexts created with configured bindings are read-only. For example, if an EJB home binding is configured with the name `some/compound/name/ejbHome`, the intermediate contexts `some`, `some/compound`, and `some/compound/name` will be created as read-only contexts. You cannot add, update, or remove any read-only bindings.

The configured binding name cannot conflict with existing bindings. However, configured bindings can use the same intermediate context names. Therefore, a configured binding with the name `some/compound/name2/ejbHome2` does not conflict with the previous example name.

## Configured binding types

Types of objects that you can bind follow:

### **EJB: EJB home installed in some server in the cell**

The following data is required to configure an EJB home binding:

- JNDI name of the EJB server or server cluster where the enterprise bean is deployed
- Target root for the configured binding (scope)
- The name of the configured binding, relative to the target root.

This type of binding is of special significance because you can use it to provide interoperability with WebSphere Application Server v3.5.x and v4.0.x JNDI clients. The default initial context for these earlier clients is the cell persistent root, which is different from the initial context of the server root for WebSphere Application Server V5 JNDI clients. If you migrate an application to the current release, you can configure an EJB binding at the cell scope so that the lookup names for the enterprise bean do not change for clients still running in a earlier WebSphere Application Server version.

A cell-scoped EJB binding is also useful for creating a fixed lookup name for an enterprise bean so that the qualified name is not dependent on the topology.

### **CORBA: CORBA object available from some CosNaming name server**

You can identify any CORBA object bound into some INS compliant CosNaming server with a `corbaname` URL. The referenced object does not have to be available until the binding is actually referenced by some application.

The following data is required in order to configure a CORBA object binding:

- The `corbaname` URL of the CORBA object
- An indicator if the bound object is a context or leaf node object (to set the correct CORBA binding type of context or object).
- Target root for the configured binding
- The name of the configured binding, relative to the target root.

### **Indirect: Any object bound in WebSphere Application Server name space accessible with JNDI**

Besides CORBA objects, this includes `javax.naming.Referenceable`, `javax.naming.Reference`, and `java.io.Serializable` objects. The target object itself is not bound to the name space. Only the information required to look up the object is bound. Therefore, the referenced name server does not have to be running until the binding is actually referenced by some application. The following data is required in order to configure an indirect JNDI lookup binding:

- JNDI provider URL of name server where object resides
- JNDI lookup name of object
- Target root for the configured binding (scope)
- The name of the configured binding, relative to the target root.

A cell-scoped indirect binding is useful when creating a fixed lookup name for a resource so that the qualified name is not dependent on the topology. You can also achieve this topology by widening the scope of the resource definition.

**Note:** WebSphere Application Server v3.5.x clients cannot access this type of binding .

**String: String constant**

You can configure a binding of a string constant. The following data is required to configure a string constant binding:

- String constant value
- Target root for the configured binding (scope)
- The name of the configured binding, relative to the target root.

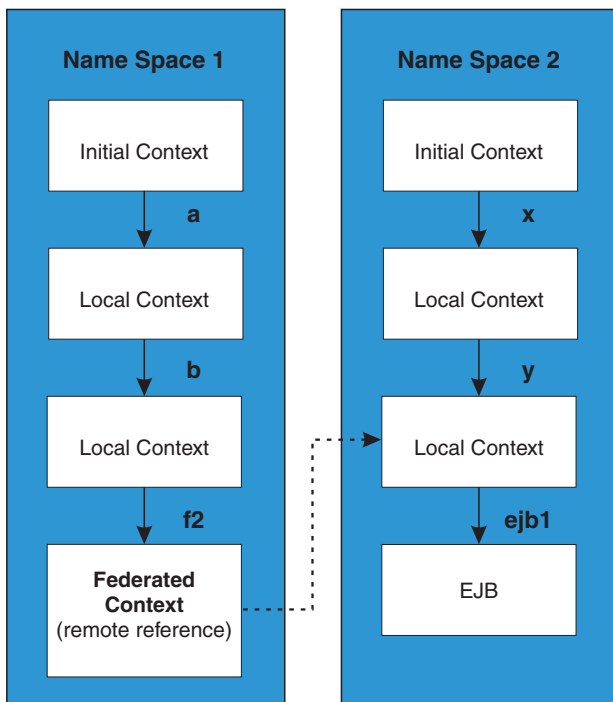
---

## Name space federation

Federating name spaces involves binding contexts from one name space into another name space.

For example, assume that a name space, Name Space 1, contains a context under the name a/b. Also assume that a second name space, Name Space 2, contains a context under the name x/y. (See the following illustration.) If context x/y in Name Space 2 is bound into context a/b in Name Space 1 under the name f2, the two name spaces are federated. Binding f2 is a federated binding because the context associated with that binding comes from another name space. From Name Space 1, a lookup of the name a/b/f2 returns the context bound under the name x/y in Name Space 2. Furthermore, if context x/y contains an Enterprise JavaBeans (EJB) home bound under the name ejb1, the EJB home could be looked up from Name Space 1 with the lookup name a/b/f2/ejb1. Notice that the name crosses name spaces. This fact is transparent to the naming client.

### Federated Name Spaces



In a WebSphere Application Server name space, you can create federated bindings with the following restrictions:

- Federation is limited to CosNaming name servers. A WebSphere Application Server name server is a Common Object Request Broker Architecture (CORBA) CosNaming implementation. You can create federated bindings to other CosNaming contexts. You cannot, for example, bind contexts from an LDAP name server implementation.



- If you use JNDI to federate the name space, you must use WebSphere Application Server's initial context factory to obtain the reference to the federated context. If you use some other initial context factory implementation, you either may not be able to create the binding, or the level of transparency may be reduced.
- A federated binding to a non-WebSphere Application Server naming context has the following functional limitations:
  - JNDI operations are restricted to the use of CORBA objects. For example, you can look up EJB homes, but you cannot look up non-CORBA objects such as data sources.
  - JNDI caching is not supported for non-WebSphere Application Server name spaces. This restriction affects the performance of lookup operations only.
- Do not federate two WebSphere Application Server stand-alone server name spaces. Incorrect behavior may result. If you want to federate WebSphere Application Server name spaces, you should use servers running under the Network Deployment or Enterprise packages of WebSphere Application Server.

---

## Name space bindings

Administrators can add name bindings to the name space through the configuration. Name servers add these configured bindings to the name space view by reading the configuration data for the bindings. Configuring bindings is an alternative to creating the bindings from a program.

Configured bindings are created each time a server starts, even when the binding is created in a transient partition of the name space. One major use of configured bindings is to provide interoperability with JNDI clients running on previous versions of the WebSphere Application Server.

There are four different kinds of bindings that you can configure:

- Enterprise JavaBeans (EJB)
- CORBA object
- Indirect Lookup
- String

---

## Configuring and viewing name space bindings

To view or configure an EJB, CORBA, Indirect lookup or string name space binding, complete the following:

1. Open the Administrative console.
2. Click **Environment**.
3. Click **Manage Name Space Bindings**.
4. Select the desired scope by entering in a node name for node-scoped bindings, or a node name and server name for server-scoped bindings, and click **Apply**.
5. To create a new binding, click **New** and follow the instructions. To edit a previously created binding, click the binding you want to edit and proceed to the next step.
6. Edit the Binding identifier, the Name in name space, and the String value fields as desired.

**Note:** All of these fields are required.

7. Click **Finish** to register the changes.

## String binding settings

Use this page to configure a new string binding or to view or edit an existing string binding.

To view this administrative console page, click **Environment** > **Naming** > **Name Space Bindings** > *string\_namespace\_binding*.

## Scope

Shows the scope of the configured binding. This value indicates the configuration location for the `namebindings.xml` file. This field is for information purposes only and cannot be updated.

If the configured binding is cell-scoped, the starting context is the cell persistent root context. If the configured binding is node-scoped, the starting context is the node persistent root context. If the configured binding is server-scoped, the starting context is the server's server root context.

## Binding Type

Shows the type of binding configured. Possible choices are String, EJB, CORBA, and Indirect. This field is for information purposes only and cannot be updated.

## Binding Identifier

Specifies the name that uniquely identifies this configured binding.

## Name in Name Space

Specifies the name used for this binding in the name space. This name can be a simple or compound name depending on the portion of the name space where this binding is configured.

## String Value

Specifies the string to be bound into the name space.

## CORBA object binding settings

Use this page to configure a new name binding of a CORBA object binding, or to view or edit an existing CORBA object binding.

To view this administrative console page, click **Environment > Naming > Name Space Bindings > CORBA\_namespace\_binding**.

## Scope

Shows the scope of the configured binding. This value indicates the configuration location for the `namebindings.xml` file. This field is for information purposes only and cannot be updated.

If the configured binding is cell-scoped, the starting context is the cell persistent root context. If the configured binding is node-scoped, the starting context is the node persistent root context. If the configured binding is server-scoped, the starting context is the server's server root context.

## Binding Type

Shows the type of binding configured. Possible choices are String, EJB, CORBA, and Indirect. This field is for information purposes only and cannot be updated.

## Binding Identifier

Specifies the name that uniquely identifies this configured binding.

## Name in Name Space

Specifies the name used for this binding in the name space. This name can be a simple or compound name depending on the portion of the name space where this binding is configured.

## Corbaname URL

Specifies the CORBA name URL string identifying where the object is bound in a CosNaming server.

## Federated Context

Specifies whether the target is a CosNaming context (true) or a leaf node object (false).

Value

Result

<b>true</b>	The target object is bound with a context CORBA binding type. If the corbaname URL does not resolve to a NamingContext, an error occurs when the binding is first used (which is when the URL is first resolved).
<b>false</b>	The target object is bound with an object CORBA binding type.

## Indirect lookup binding settings

Use this page to configure a new indirect lookup name binding, or to view or edit an existing indirect lookup binding.

To view this administrative console page, click **Environment > Naming > Name Space Bindings > indirect\_lookup\_namespace\_binding**.

### Scope

Shows the scope of the configured binding. This value indicates the configuration location for the namebindings.xml file. This field is for information purposes only and cannot be updated.

If the configured binding is cell-scoped, the starting context is the cell persistent root context. If the configured binding is node-scoped, the starting context is the node persistent root context. If the configured binding is server-scoped, the starting context is the server's server root context.

### Binding Type

Shows the type of binding configured. Possible choices are String, EJB, CORBA, and Indirect. This field is for information purposes only and cannot be updated.

### Binding Identifier

Specifies the name that uniquely identifies this configured binding.

### Name in Name Space

Specifies the name used for this binding in the name space. This name can be a simple or compound name depending on the portion of the name space where this binding is configured.

### Provider URL

Specifies the provider URL string needed to obtain a JNDI initial context.

### JNDI Name

Specifies the name used to look up the target object from the initial context.

## EJB binding settings

Use this page to configure a new EJB binding, or to view or edit an existing EJB binding.

To view this administrative console page, click **Environment > Naming > Name Space Bindings > EJB\_namespace\_binding**.

### Scope

Shows the scope of the configured binding. This value indicates the configuration location for the namebindings.xml file. This field is for information purposes only and cannot be updated.

If the configured binding is cell-scoped, the starting context is the cell persistent root context. If the configured binding is node-scoped, the starting context is the node persistent root context. If the configured binding is server-scoped, the starting context is the server's server root context.

## Binding Type

Shows the type of binding configured. Possible choices are String, EJB, CORBA, and Indirect. This field is for information purposes only and cannot be updated.

## Binding Identifier

Specifies the name that uniquely identifies this configured binding.

## Name in Name Space

Specifies the name used for this binding in the name space. This name can be a simple or compound name depending on the portion of the name space where this binding is configured.

## Enterprise Bean Location

Specifies whether the enterprise bean is running in a server cluster or a single server. If Single Server is specified, type the node name.

## Server

Specifies the name of the cluster or non-clustered server in which the enterprise bean is configured.

## JNDI Name

Specifies the JNDI name of the deployed enterprise bean (the bean's JNDI name that is in the enterprise bean bindings--not the java:comp name)

## Name space binding collection

Use this page to configure a name binding of an EJB, a CORBA CosNaming NamingContext, a CORBA leaf node object, an object that you can look up using JNDI, or a constant string value.

Binding information for configured bindings is stored in the configuration and applied upon startup of the name server for each server within the scope of the binding.

To view the Manage Name Space Bindings Settings page, click **Environment > Naming > Name Space Bindings**.

Click the check boxes to select one or more of the users in your collection. Use the buttons to control the selected users.

## Name

Shows the names given to uniquely identify these configured bindings.

## Scope

Shows the scope of the configured binding. This value indicates the configuration location for the namebindings.xml file. This field is for information purposes only and cannot be updated.

If the configured binding is cell-scoped, the starting context is the cell persistent root context. If the configured binding is node-scoped, the starting context is the node persistent root context. If the configured binding is server-scoped, the starting context is the server's server root context.

## Binding Type

Shows the type of binding configured. Valid values are String, EJB, CORBA, and Indirect. This field is for information purposes only and cannot be updated.

---

## Configuring name servers

To configure a name server, complete the following:

1. Open the administrative console.
2. Click **Servers**.
3. Click **Application Servers**.
4. Click the application server you want to configure.
5. Click **Server Components**.
6. Click **Name Server**.
7. Edit the fields as desired.

**Note:** All of these fields are mandatory.

8. To make other changes, click **Custom Properties**.
9. Click **OK** to register your changes.

### Name server settings

Use this page to configure Naming Service Provider settings for the application server.

To view this administrative console page, click one of the following paths:

- **Servers > Application Servers > *server\_name* > Server Components > Name Server**
- **Servers > JMS Servers > *server\_name* > Server Components > Name Server**

#### Name

Specifies the display name for the server.

<b>Data type</b>	String
------------------	--------

#### Initial State

Specifies the execution state. The options are: *Started* and *Stopped*.

<b>Data type</b>	String
<b>Default</b>	Started

---

## Troubleshooting name space problems

Many naming problems can be avoided by fully understanding the key underlying concepts of WebSphere Application Server naming.

1. Review the key concepts of WebSphere Application Server naming, especially Name space logical view and Lookup names support in deployment descriptors and thin clients.
2. Review the programming examples that are included in the sections explaining the JNDI and CosNaming interfaces.
3. Read "Naming services component troubleshooting tips" in the Information Center for additional general information.
4. Read "Cannot look up an object hosted by WebSphere Application Server from a servlet, JSP file, or other client" in the Information Center if you encounter the described problem.

## dumpNameSpace tool

You can use the dumpNameSpace tool to dump the contents of a name space accessed through a name server. When you invoke the dumpNameSpace tool, the naming service must be active. The dumpNameSpace tool cannot dump name spaces local to the server process, such as those with java: and local: URL schemes. The local: name space contains references to enterprise beans with local interfaces. Use the name space dump utility for java:, local: and server name spaces to dump java: and local: name spaces.

Note that the server root context for the server at the specified host and port is dumped (unless a non-default starting context which precludes it is specified). The server root contexts for other servers are not dumped.

If you run the dumpNameSpace tool, a login prompt is displayed. If you cancel the login prompt, the dumpNameSpace tool continues outbound with an "UNAUTHENTICATED" credential. Thus, by default, an "UNAUTHENTICATED" credential is used that is equivalent to the "Everyone" access authorization policy. You can modify this default setting by changing the value for the com.ibm.CSI.performClientAuthenticationRequired property to true in the install\_dir/properties/sas.client.props file. When you change this property to true, rerun the dumpNameSpace tool, and cancel the login prompt, the authorization fails and the command will not continue outbound.

Command line invocation descriptions of the dumpNameSpace tool follow. This section includes sample output.

You can also access this tool through its program interface. Refer to the class com.ibm.websphere.naming.DumpNameSpace in the WebSphere Application Server API documentation. To invoke the tool through the command line, enter the following command from the WebSphere/AppServer/bin directory:

Platform	Command
UNIX	dumpNameSpace.sh <i>[-keyword value]...</i>
Windows NT	dumpNameSpace <i>[-keyword value]...</i>

### Parameters

The keywords and associated values for the dumpNameSpace utility follow:

**-host** *myhost.austin.ibm.com*

Indicates the bootstrap host or the WebSphere Application Server host whose name space you want to dump. The value defaults to **localhost**.

**-port** *nnn*

Indicates the bootstrap port which, if not specified, defaults to **2809**.

**-root** {*cell | server | node | host | legacy | tree | default*}

Indicates the root context to use as the initial context for the dump. The applicable root options and default root context depend on the type of name server from which the dump is being obtained. This information is provided in the following tables.

For WebSphere Application Servers V5 or later:

cell	DumpNameSpace default. Dump the tree starting at the cell root context.
server	Dump the tree starting at the server root context.
node	Dump the tree starting at the node root context. (Synonymous with host.)

For WebSphere Application Servers v4.0 or later:

legacy	DumpNameSpace default. Dump the tree starting at the legacy root context.
host	Dump the tree starting at the bootstrap host root context. (Synonymous with node.)
tree	Dump the tree starting at the tree root context.

For all WebSphere Application Servers and other name servers:

default	Dump the tree starting at the initial context which JNDI returns by default for that server type. This is the only -root choice that is compatible with WebSphere Application Servers prior to v4.0 and with non-WebSphere Application Server name servers.
---------	-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

**-url** *some provider URL*

Indicates the value for the java.naming.provider.url property used to get the initial JNDI context. This option can be used in place of the -host, -port, and -root options. If the -url option is specified, the -host, -port, and -root options are ignored.

**-factory** *com.ibm.websphere.naming.WsnInitialContextFactory*

Indicates the initial context factory to be used to get the JNDI initial context. The value defaults to: *com.ibm.websphere.naming.WsnInitialContextFactory* The default value generally does not need to be changed.

**-startAt** *some/subcontext/in/the/tree*

Indicates the path from the bootstrap host's root context to the top level context where the dump should begin. The utility recursively dumps subcontexts below this point. It defaults to an empty string, that is, the bootstrap host root context.

**-format{jndi | ins}**

Option	Description
jndi	The default. Displays name components as atomic strings.
ins	Shows name components parsed per INS rules (id.kind).

**-report {short | long}**

Option	Description
short	The default. Dumps the binding name and bound object type. This output is also provided by JNDI Context.list().
long	<p>Dumps the binding name, bound object type, local object type, and string representation of the local object (that is, the IORs, string values, and other values that are printed).</p> <p>For objects of user-defined classes to display correctly with the long report option, it may be necessary to add their containing directories to the list of directories searched. Set the environment variable <b>WAS_USER_DIRS</b>. The value can include one or more directories, as for example:</p> <p><b>Platform</b></p> <p><b>Command</b></p> <p><b>UNIX</b> WAS_USER_DIRS=/usr/classdir1:/usr/classdir2 export WAS_USER_DIRS</p> <p><b>Windows NT</b> set WAS_USER_DIRS=c:\classdir1;d:\classdir2</p> <p>All zip, jar, and class files in the specified directories can then be resolved by the class loader when running <b>dumpNameSpace</b>.</p>



**-traceString** *"some.package.name.to.trace.\*=all=enabled"*

Represents the trace string with the same format as that generated by the servers. The output is sent to the file, DumpNameSpaceTrace.out.

## Example: Invoking the name space dump utility

It is often helpful to view a dump of the name space to understand why a naming operation is failing. You can invoke the name space dump utility from the command line or from a program. Examples of each option follow.

### Invoking name space dump utility from a command line

Invoke the name space dump utility from the command line by entering the following command:

```
dumpNameSpace -host myhost.mycompany.com -port 901
```

**OR**

```
dumpNameSpace -url corbaloc:iiop:myhost.mycompany.com:901
```

There are several command line options to choose from. For detailed help, enter the following command:

```
dumpNameSpace -help
```

### Invoking name space dump utility from a Java program

You can dump name spaces from a program with the `com.ibm.websphere.naming.DumpNameSpace` API. Refer to the WebSphere Application Server API documentation for details on the `DumpNameSpace` program interface

The following example illustrates how to invoke the name space dump utility from a Java program:

```
{
 ...
 import javax.naming.Context;
 import javax.naming.InitialContext;
 import com.ibm.websphere.naming.DumpNameSpace;
 ...
 java.io.PrintStream filePrintStream = ...
 Context ctx = new InitialContext();
 // Starting context for dump
 ctx = (Context) ctx.lookup("cell/nodes/node1/servers/server1");
 DumpNameSpace dumpUtil =
 new DumpNameSpace(filePrintStream, DumpNameSpace.SHORT);
 dumpUtil.generateDump(ctx);
 ...
}
```

## Name space dump utility for java:, local: and server name spaces

Sometimes it is helpful to dump the `java:` name space for a J2EE application. You cannot use the `dumpNameSpace` command line utility for this purpose because the application's `java:` name space is accessible only by that J2EE application. From the WebSphere Application Server scripting tool, you can invoke a `NameServer` MBean to dump the `java:` name space for any J2EE application running in that same server process.

There is another name space local to server process which you cannot dump with the `dumpNameSpace` command line utility. This name space has the URL scheme of `local:` and is used by the container to bind objects locally instead of through the name server. The `local:` name space contains references to enterprise beans with local interfaces. There is only one `local:` name space in a server process. You can dump the `local:` name space by invoking the `NameServer` MBean associated with that server process.

## Name space dump options

Name space dump options are specified in the MBean invocation as a parameter in character string format. The option descriptions follow.

### **-startAt** *some/subcontext/in/the/tree*

Indicates the path from the name space root context to the top level context where the dump should begin. The utility recursively dumps subcontexts below this point. It defaults to an empty string, that is, the root context.

### **-report {short | long}**

Option	Description
short	The default. Dumps the binding name and bound object type. This output is also provided by JNDI Context.list().
long	Dumps the binding name, bound object type, local object type, and string representation of the local object (that is, the IORs, string values, and other values that are printed).

### **-root {tree | host | legacy | cell | node | server | default}**

Specify the root context of where the dump should start. The default value for `-root` is `cell`. This option is only valid for server name space dumps.

Option	Description
tree	Dump the tree starting at the tree root context.
host	Dump the tree starting at the server host root context (synonymous with "node").
legacy	Dump the tree starting at the legacy root context.
cell	Dump the tree starting at the cell root context. This is the default option.
node	Dump the tree starting at the node root context (synonymous with "host").
server	Dump the tree starting at the server root context. This is <code>-root</code> default.
default	Dump the tree starting at the initial context which JNDI returns by default for that server type.

### **-format {jndi | ins}**

Specify the format to display name component as atomic strings or parsed according to INS rules (id.kind). This option is only valid for server name space dumps.

Option	Description
jndi	Display name components as atomic strings. This is <code>-format</code> default.
ins	Display name components parsed according to INS rules (id.kind).

## NameServer MBean invocation

1. Enter the WebSphere Application Server scripting command prompt.

Invoke a method on a NameServer MBean by using the WebSphere Application Server scripting tool. Enter the scripting command prompt by typing the following command:

Platform	Command
UNIX	wsadmin.sh
Windows NT	wsadmin

Use the `-help` option for help on using the `wsadmin` command.

2. Select the NameServer MBean instance to invoke.

Execute the following script commands to select the NameServer instance you want to invoke. For example,

```
set mbean [$AdminControl completeObjectName WebSphere:*,type=NameServer,cell=
 cellName,node=nodeName,process=serverName]
```

where *cellName*, *nodeName*, and *serverName* are the names of the cell, node, and server for the MBean you want to invoke. The specified server must be running before you can invoke a method on the MBean.

You can see a list of all NameServer MBeans current running by issuing the following query:

```
$AdminControl queryNames {*:*,type=NameServer}
```

3. Invoke the NameServer MBean.

#### **java: name space**

Dump a java: name space by invoking the `dumpJavaNameSpace` method on the NameServer MBean. Since each server application has its own java: name space, the application must be specified on the method invocation. An application is identified by the application name, module name, and component name. The method syntax follows:

```
$AdminControl invoke $mbean dumpJavaNameSpace {{appName}{modName}{compName}{opts}}
```

where *appName* is the application name, *modName* is the module name, and *compName* is the component name of the java: name space you want to dump. The value for *opts* is the list of name space dump options described earlier in this section. The list can be empty.

#### **local: name space**

Dump a java: name space by invoking the `dumpLocalNameSpace` method on the NameServer MBean. Since there is only one local: name space in a server process, you have to specify the name space dump options only.

```
$AdminControl invoke $mbean dumpLocalNameSpace {{opts}}
```

where *opts* is the list of name space dump options described earlier in this section. The list can be empty.

#### **Server name space**

Dump a server name space by invoking the `dumpServerNameSpace` method on an application server's NameServer MBean. This provides an alternative way to dump the name space on an application server, much like the `dumpNameSpace` command line utility.

```
$AdminControl invoke $mbean dumpServerNameSpace {{opts}}
```

where *opts* is the list of name space dump options described earlier in this section. The list can be empty.

### **Name space dump output**

Name space dump output is sent to the console. It is also written to the file `DumpNameSpace.log`, in the server's log directory.

## **Example: Invoking the name space dump utility for java: and local: name spaces**

It is often helpful to view the dump of a java: or local: name space to understand why a naming operation is failing. The NameServer MBean running in the application's server process can be invoked from the WebSphere Application Server scripting tool to generate a dump of these name spaces. Examples of NameServer MBean calls to generate dumps of java: and local: name spaces follow.

## Dumping a java: name space

Assume you want to dump the java: name space of an application component running in server server1 on node node1 of the cell MyCell. The application name is AcctApp in module AcctApp.war, and the component name is Acct Servlet. The following script commands generate a long format dump of the application's java: name space of that application:

```
set mbean [$AdminControl completeObjectName WebSphere:*,type=NameServer,cell=MyCell,node=node1,process=server1]
$AdminControl invoke $mbean dumpJavaNameSpace {{AcctApp}{AcctApp.war}{Acct Servlet}{-report long}}
```

## Dumping a local: name space

Assume you want to dump the local: name space for the server server1 on node node1 of cell MyCell. The following script commands will generate a short format dump of that server's local name space:

```
set mbean [$AdminControl completeObjectName WebSphere:*type=NameServer,cell=MyCell,node=node1,process=server1]
$AdminControl invoke $mbean dumpLocalNameSpace {{-report short}}
```

## Name space dump sample output

Name space dump output looks like the following example, which is the **SHORT** dump format:

```
Getting the initial context
Getting the starting context
```

```
=====
Name Space Dump
```

```
 Provider URL: corbaloc:iiop:localhost:9810
 Context factory: com.ibm.websphere.naming.WsnInitialContextFactory
 Requested root context: cell
 Starting context: (top)=outpostNetwork
 Formatting rules: jndi
 Time of dump: Mon Sep 16 18:35:03 CDT 2002
=====
```

```
=====
Beginning of Name Space Dump
=====
```

```
 1 (top)
 2 (top)/domain javax.naming.Context
 2 Linked to context: outpostNetwork
 3 (top)/cells javax.naming.Context
 4 (top)/clusters javax.naming.Context
 5 (top)/clusters/Cluster1 javax.naming.Context
 6 (top)/cellname java.lang.String
 7 (top)/cell javax.naming.Context
 7 Linked to context: outpostNetwork
 8 (top)/deploymentManager javax.naming.Context
 8 Linked to URL: corbaloc::outpost:9809/NameServiceServerRoot
 9 (top)/nodes javax.naming.Context
10 (top)/nodes/will2 javax.naming.Context
11 (top)/nodes/will2/persistent javax.naming.Context
12 (top)/nodes/will2/persistent/SomeObject SomeClass
13 (top)/nodes/will2/nodename java.lang.String
14 (top)/nodes/will2/domain javax.naming.Context
14 Linked to context: outpostNetwork
15 (top)/nodes/will2/cell javax.naming.Context
15 Linked to context: outpostNetwork
16 (top)/nodes/will2/servers javax.naming.Context
17 (top)/nodes/will2/servers/server1 javax.naming.Context
18 (top)/nodes/will2/servers/will2 javax.naming.Context
19 (top)/nodes/will2/servers/member2 javax.naming.Context
20 (top)/nodes/will2/node javax.naming.Context
20 Linked to context: outpostNetwork/nodes/will2
21 (top)/nodes/will2/nodeAgent javax.naming.Context
```

22 (top)/nodes/outpost javax.naming.Context  
 23 (top)/nodes/outpost/node javax.naming.Context  
 23 Linked to context: outpostNetwork/nodes/outpost  
 24 (top)/nodes/outpost/nodeAgent javax.naming.Context  
 24 Linked to URL: corbaloc::outpost:2809/NameServiceServerRoot  
 25 (top)/nodes/outpost/persistent javax.naming.Context  
 26 (top)/nodes/outpost/nodename java.lang.String  
 27 (top)/nodes/outpost/domain javax.naming.Context  
 27 Linked to context: outpostNetwork  
 28 (top)/nodes/outpost/servers javax.naming.Context  
 29 (top)/nodes/outpost/servers/server1 javax.naming.Context  
 30 (top)/nodes/outpost/servers/server1/url javax.naming.Context  
 31 (top)/nodes/outpost/servers/server1/url/CatalogDAOURL  
 31 java.net.URL  
 32 (top)/nodes/outpost/servers/server1/mail javax.naming.Context  
 33 (top)/nodes/outpost/servers/server1/mail/PlantsByWebSphere  
 33 javax.mail.Session  
 34 (top)/nodes/outpost/servers/server1/TransactionFactory  
 34 com.ibm.ejs.jts.jts.ControlSet\$LocalFactory  
 35 (top)/nodes/outpost/servers/server1/servername java.lang.String  
 36 (top)/nodes/outpost/servers/server1/WSsamples javax.naming.Context  
 37 (top)/nodes/outpost/servers/server1/WSsamples/TechSampDatasource  
 37 TechSamp  
 38 (top)/nodes/outpost/servers/server1/thisNode javax.naming.Context  
 38 Linked to context: outpostNetwork/nodes/outpost  
 39 (top)/nodes/outpost/servers/server1/cell javax.naming.Context  
 39 Linked to context: outpostNetwork  
 40 (top)/nodes/outpost/servers/server1/eis javax.naming.Context  
 41 (top)/nodes/outpost/servers/server1/eis/DefaultDatasource\_CMP  
 41 Default\_CF  
 42 (top)/nodes/outpost/servers/server1/eis/WSsamples javax.naming.Context  
 43 (top)/nodes/outpost/servers/server1/eis/WSsamples/TechSampDatasource\_CMP  
 43 TechSamp\_CF  
 44 (top)/nodes/outpost/servers/server1/eis/jdbc javax.naming.Context  
 45 (top)/nodes/outpost/servers/server1/eis/jdbc/PlantsByWebSphereDataSource\_CMP  
 45 PLANTSDB\_CF  
 46 (top)/nodes/outpost/servers/server1/eis/jdbc/petstore  
 46 javax.naming.Context  
 47 (top)/nodes/outpost/servers/server1/eis/jdbc/petstore/PetStoreDB\_CMP  
 47 PetStore\_CF  
 48 (top)/nodes/outpost/servers/server1/eis/jdbc/CatalogDB\_CMP  
 48 Catalog\_CF  
 49 (top)/nodes/outpost/servers/server1/jta javax.naming.Context  
 50 (top)/nodes/outpost/servers/server1/jta/usertransaction  
 50 java.lang.Object  
 51 (top)/nodes/outpost/servers/server1/DefaultDatasource  
 51 Default Datasource  
 52 (top)/nodes/outpost/servers/server1/jdbc javax.naming.Context  
 53 (top)/nodes/outpost/servers/server1/jdbc/CatalogDB CatalogDB  
 54 (top)/nodes/outpost/servers/server1/jdbc/petstore javax.naming.Context  
 55 (top)/nodes/outpost/servers/server1/jdbc/petstore/PetStoreDB  
 55 PetStoreDB  
 56 (top)/nodes/outpost/servers/server1/jdbc/PlantsByWebSphereDataSource  
 56 PLANTSDB  
 57 (top)/nodes/outpost/servers/outpost javax.naming.Context  
 57 Linked to URL: corbaloc::outpost:2809/NameServiceServerRoot  
 58 (top)/nodes/outpost/servers/member1 javax.naming.Context  
 59 (top)/nodes/outpost/cell javax.naming.Context  
 59 Linked to context: outpostNetwork  
 60 (top)/nodes/outpostManager javax.naming.Context  
 61 (top)/nodes/outpostManager/domain javax.naming.Context  
 61 Linked to context: outpostNetwork  
 62 (top)/nodes/outpostManager/cell javax.naming.Context  
 62 Linked to context: outpostNetwork  
 63 (top)/nodes/outpostManager/servers javax.naming.Context  
 64 (top)/nodes/outpostManager/servers/dmgr javax.naming.Context  
 64 Linked to URL: corbaloc::outpost:9809/NameServiceServerRoot

```

65 (top)/nodes/outpostManager/node javax.naming.Context
65 Linked to context: outpostNetwork/nodes/outpostManager
66 (top)/nodes/outpostManager/nodename java.lang.String
67 (top)/persistent javax.naming.Context
68 (top)/persistent/cell javax.naming.Context
68 Linked to context: outpostNetwork
69 (top)/legacyRoot javax.naming.Context
69 Linked to context: outpostNetwork/persistent
70 (top)/persistent/AnotherObject AnotherClass

```

```

=====
End of Name Space Dump
=====

```

---

## Naming and directories: Resources for learning

Use the following links to find relevant supplemental information about naming and directories. The information resides on IBM and non-IBM Internet sites, whose sponsors control the technical accuracy of the information.

These links are provided for convenience. Often, the information is not specific to the IBM WebSphere Application Server product, but is useful all or in part for understanding the product. When possible, links are provided to technical papers and Redbooks that supplement the broad coverage of the release documentation with in-depth examinations of particular product areas.

### Programming instructions and examples

- Naming in WebSphere Application Server V5: Impact on Migration and Interoperability

### Programming specifications

- Java Naming and Directory Interface™ 1.2.1 Specification
- OMG CosNaming Interoperable Naming Specification





---

## Chapter 27. Using the dynamic cache service to improve performance

The dynamic cache service works within an application server Java virtual machine (JVM), intercepting calls to cacheable objects. For example, it intercepts calls through a servlet service() method or a command execute() method, and either stores the output of the object to, or serves the content of the object from the dynamic cache.

WebSphere Application Server, Version 4.0, supported the configuration of dynamic servlet caching through the servletcache.xml file. To utilize the new and improved functionality of the dynamic cache service, configure your cache policy using the cachespec.xml format. See “Configuring cacheable objects with the cachespec.xml file” on page 1354 for more information.

The dynamic caching documentation provides you with the following tasks to enable and configure the dynamic cache service, as well as advanced features, such as controlling external caches and building user-defined drop-in components to customize the cache operation.

1. Enable the dynamic cache service globally.
2. Configure servlet caching.
3. Configure Edge Side Include (ESI) caching.
4. Configure command caching.
5. Configure Web services caching.
6. Troubleshoot any problems with the dynamic cache service.

To use the DistributedMap interface for the dynamic cache, see Using the DistributedMap interface for the dynamic cache.

---

### Dynamic cache

Caching the output of servlets, commands and JavaServer Pages (JSP) files, improves application performance. WebSphere Application Server consolidates several caching activities, including servlets, Web services, and WebSphere commands into one service called the *dynamic cache*. These caching activities work together to improve application performance, and share many configuration parameters, which are set in the dynamic cache service of an application server.

You can use the dynamic cache to improve the performance of servlet and JSP files by serving requests from an in-memory cache. Cache entries contain servlet output, results of servlet execution, and metadata.

---

### Configuring cache replication

Cache replication leverages the WebSphere internal replication service that is also leveraged for HttpSession memory-to-memory replication for failover purposes. Hence, a replication domain with at least one replicator entry needs to exist in order to replicate the data. The dynamic cache service, in essence, connects to the replicator. See more information in the topic referring to managing internal replication, “Replicating data,” in the Information Center.

To configure cache replication and its features:

1. Click **Servers > Application Servers** in the administrative console navigation tree.
2. Click *server*.
3. Click **Dynamic Cache Service**.
4. Select the **Enable cache replication** check box in the **Cache replication** field.

To manage **batch update** or **PUSH-PULL - PUSH/PULL**, repeat steps 1-4, then click the **Enable cache replication** link to the right of the **Enable cache replication** check box. Batch update interval is set under **push frequency**. PUSH-PULL-PUSH/PULL is set through the **runtime mode**.

You can also select which replication domain and initial replicator entry the dynamic cache will utilize (either those managed within the cell or across the cell).

## Cache replication

Data is generated one time and copied or replicated to other servers in the cluster, saving execution time and resources. Caching in a cluster has additional concerns. In particular, the same data can be required and generated in multiple places. Also, the permission the resources need to generate the cached data can be restricted, preventing access to the data.

Cache replication addresses these concerns by generating the data one time and copying or replicating it to the other servers in the cluster. It also aids in cache consistency, in that cache entries that are not needed are removed or replaced.

The data replication configuration can exist as part of the Web container dynamic cache configuration accessible through the administrative console, or on a per cache entry basis through the `cachespec.xml` file. With the `cachespec.xml` file, you can configure cache replication at the Web container level, but disable it for a specific cache entry.

Cache replication can take on three forms:

- **PUSH** - Send out new entries, both ID and data, and updates to those entries.
- **PULL** - Requests data from other servers in the cluster when that data is not locally present. This mode of replication is not recommended.
- **PUSH/PULL** - Sends out IDs for new entries, then, only requests from other servers in the cluster entries for IDs previously broadcast. The dynamic cache always sends out cache entry invalidations.

The dynamic cache provides a batch update option. Specifically, for PUSH or PUSH/PULL, the dynamic cache broadcasts the update asynchronously, based on a timed interval rather than sending them immediately upon inception. Invalidators are sent immediately. Distribution of invalidations prevents stale data from residing in a cluster.

## Internal messaging configuration settings

Use this page to set advanced configurations for Memory to Memory session replication.

To view this administrative console page, click **Servers > Application Servers > server > Dynamic Cache Service > Cache replication > Enable cache replication** .

The advanced replication settings include fields for choosing the initial replicator entry that connects to the replicator domains. As an alternative, you can specify the IP addresses and ports (of the form `address:port`) for connection to replicators outside of the cell that the server is administered under. By default, if a replicator is defined on the server you are configuring, that server is the one chosen for cache replication. Select the advanced properties only if you want to deviate from the default setting.

### Internal messaging server

Specifies a domain from which your data will be replicated. Depending on the domain you choose to replicate the data, you can choose any of the replicators defined under that domain. You can use the default domain or choose one from the drop down window.

### Runtime mode

Specifies the global sharing policy for this server.

The following settings are available:

- **Both push and pull** sends the cache ID of newly updated content to other servers in the replication domain. Then, if one of the other servers requests the content, and that server has the ID of the cache entry for the previously updated content, it will retrieve the content from the publishing server. On the other hand, if a request is made for an ID which has not been previously published, the server assumes it does not exist in the cluster and creates a new entry.
- **Push only** sends the cache ID and cache content of new content to all other servers in the replication domain.
- The sharing policy of **Not Shared** results in the cache ID and cache content not being shared with other servers in the replication domain.

The default setting for a non-clustered environment is **Not Shared**. When enabling replication, the default value is **Push only**.

### Push frequency

Specifies the time in seconds to wait before pushing new or modified cache entries to other servers. A value of 0 (zero) means send immediately.

Setting this property to a value greater than 0 (zero) causes a "batch" push of all cache entries that are created or modified during the time period.

Default	0 (equivalent to immediate)
---------	-----------------------------

---

## Enabling the dynamic cache service

In order to use the dynamic cache service, you must first enable it.

1. Open the administrative console.
2. Click **Servers > Application Servers** in the administrative console navigation tree.
3. Click a server.
4. Click **Dynamic Cache Service** under Additional Properties.
5. Select **Enable service at server startup** in the **Startup state** field.
6. Click **Apply** or **OK**.
7. Restart WebSphere Application Server.

The dynamic cache service will now cache content for requests that have cache policies configured.

## Dynamic cache service settings

Use this page to configure and manage the dynamic cache service settings.

To view this administrative console page, click **Servers > Application Servers > server > Dynamic Cache Service**.

### Startup state

Specifies whether the dynamic cache is enabled.

### Cache size

Specifies a positive integer as the value for the maximum number of entries the cache holds.

Enter the cache size value in this field between the range of 100 through 200,000.

### Default priority

Specifies the default priority for cache entries, determining how long an entry stays in a full cache.

Default	1
---------	---

Range	1 to 255
-------	----------

## Disk offload

Specifies whether disk offload is enabled.

By default, the dynamic cache maintains the number of entries configured in memory. If new entries are created while the cache is full, the priorities configured for each cache entry and a least recently used algorithm, are used to remove entries from the cache. In addition to having a cache entry removed from memory when the cache is full, you can enable disk offload to have a cache entry copied to the file system (the location is configurable). Later, if that cache entry is needed, it is moved back to memory from the file system.

## Cache replication

Specifies whether cache replication is enabled.

You can also configure advanced cache replication settings.

## Configuring servlet caching

To enable servlet caching, you must complete “Enabling the dynamic cache service” on page 1345.

1. Open the administrative console.
2. Click **Servers > Application Servers** in the console navigation tree.
3. Click a server.
4. Click **Web Container**.
5. Select **Enable servlet caching** under the Configuration tab.
6. Click **Apply** or **OK**.

Define the cache policy for your servlets by “Configuring cacheable objects with the cachespec.xml file” on page 1354.

## Servlet caching

After a servlet is invoked and generating the output to cache, a cache entry is created containing the output and the side effects of the invocation. For example, these side effects can include calls to other servlets or JavaServer Pages (JSP) files, as well as metadata about the entry, including timeout and entry priority information.

Unique entries are distinguished by an ID string generated from the `HttpServletRequest` object for each invocation of the servlet. You can then base servlet caching on:

- Request parameters and attributes the URI used to invoke the servlet
- Session information
- Other options, including cookies

Since JSP files are compiled by WebSphere Application Server into servlets, the dynamic cache function treats them the same, except in specifically documented situations.

## Configuring the dynamic cache disk offload

By default, when the number of cache entries reaches the configured limit for a given WebSphere server, eviction of cache entries occurs, allowing new entries to enter the cache service. The dynamic cache includes an alternative feature named disk offload, that copies the evicted cache entries to disk for potential future access.

To configure disk offload:

1. Open the administrative server.
2. Click **Server > Application Server** in the administrative console navigation tree.
3. Click *server*.
4. Click **Dynamic Cache Service**.
5. Click the **Enable disk offload** check box in the **Disk offload** field. You can also set the disk offload location in this field.
6. Click **Apply** or **OK**.

### Application servers must have different disk offload locations

When you have two or more application servers with servlet caching enabled and the application servers specify the same disk offload location for their caches through the dynamic cache service, the following exceptions might occur:

```
java.lang.NullPointerException
 at com.ibm.ws.cache.CacheOnDisk.readTemplate(CacheOnDisk.java:686)
 at com.ibm.ws.cache.Cache.internalInvalidateByTemplate(Cache.java:828)
```

or:

```
java.lang.NullPointerException
 at com.ibm.ws.cache.CacheOnDisk.readCacheEntry(CacheOnDisk.java:600)
 at com.ibm.ws.cache.Cache.getCacheEntry(Cache.java:341)
```

If one server is run as root and the other servers are run as nonroot, this problem could occur. For example, if server1 runs as root and server2 runs as wasuser or wasgroup, the cache files in the disk offload location might be created with root permissions. This situation causes the applications running on the nonroot servers to crash when they try to read or write to the cache.

The disk offload location must be unique for servers defined on the same node. If you have multiple servers defined on the same node, make sure the disk offload location is different for each server as defined on the **Dynamic Cache Service** panel, **Offload location** field.

## Configuring Edge Side Include caching

Edge Side Include (ESI) is configured through the `plugin-cfg.xml` file.

The Web server plug-in contains a built-in ESI processor. The ESI processor has the ability to cache whole pages, as well as fragments, providing a higher cache hit ratio. The cache implemented by the ESI processor is an in-memory cache, not a disk cache, therefore, the cache entries are not saved when the Web server is restarted.

The basic operation of the ESI processor is as follows: When a request is received by the Web server plug-in, it is sent to the ESI processor, unless the ESI processor is disabled. It is enabled by default. If a cache miss occurs, a `Surrogate-Capabilities` header is added to the request and the request is forwarded to the WebSphere Application Server. If the dynamic servlet cache is enabled in the application server, and the response is edge cacheable, the application server returns a `Surrogate-Control` header in response to the WebSphere Application Server plug-in.

The value of the `Surrogate-Control` response header contains the list of rules which are used by the ESI processor in order to generate the cache ID. The response is then stored in the ESI cache, using the cache ID as the key. For each ESI include tag in the body of the response, a new request is processed such that each nested include results in either a cache hit or another request forwarded to the application server. When all nested includes have been processed, the page is assembled and returned to the client.

The ESI processor is configurable through the WebSphere Web server plug-in configuration file `plugin-cfg.xml`. The following is an example of the beginning of this file, which illustrates the ESI configuration options.

```
<?xml version="1.0"?>
<Config>
 <Property Name="esiEnable" Value="true"/>
 <Property Name="esiMaxCacheSize" Value="1024"/>
 <Property Name="esiInvalidationMonitor" Value="false"/>
</Config>
```

The first option, `esiEnable`, can be used to disable the ESI processor by setting the value to false. ESI is enabled by default. If ESI is disabled, then the other ESI options are ignored.

The second option, `esiMaxCacheSize`, is the maximum size of the cache in 1K byte units. The default maximum size of the cache is 1 megabyte. If the cache is full, the first entry to be evicted from the cache is the entry that is closest to expiration.

The third option, `esiInvalidationMonitor`, specifies whether or not the ESI processor should receive invalidations from the application server. ESI works well when the Web servers following a threading model is used, and only one process is started. When multiple processes are started, each process caches the responses independently and the cache is not shared. This could lead to a situation where, the system's memory is fully used up by ESI processor. There are three methods by which entries are removed from the ESI cache: first, an entry's expiration timeout could fire; second, an entry may be purged to make room for newer entries; or third, the application server could send an explicit invalidation for a group of entries. In order for the third mechanism to be enabled, the `esiInvalidationMonitor` property must be set to true and the `DynaCacheEsi` application must be installed on the application server. The `DynaCacheEsi` application is located in the `installableApps` directory and is named `DynaCacheEsi.ear`. If the `ESIInvalidationMonitor` property is set to true but the `DynaCacheEsi` application is not installed, then errors will occur in the webserver plugin and the request will fail.

The ESI processor's cache can be monitored through the `CacheMonitor` application. In order for ESI processor's cache to be visible in the `CacheMonitor`, the `DynaCacheEsi` application must be installed as described above and the `ESIInvalidationMonitor` property must be set to true in the `plugin-cfg.xml` file.

When WebSphere Application Server is used to serve static data, such as images and HTML on the application server, the URLs are also cached in the ESI processor. This data has a default timeout of 300 seconds. You can change the timeout value by adding the property `com.ibm.servlet.file.esi.timeOut` to your JVM's command line parameters. The following example shows how to set a one minute timeout on static data cached in the plug-in:

```
-Dcom.ibm.servlet.file.esi.timeOut=60
```

For more information about the `plugin-cfg.xml` file see "Manually editing the plug-in configuration" in the Information Center.

For information about configuring alternate URL, see "Configuring alternate URL."

## Configuring alternate URL

Alternate URL is a method for edge caching JavaServer Pages (JSP) files and servlet responses that you can not request externally. Dynamic cache provides support to recognize the presence of an Edge Side Include (ESI) processor and to generate ESI include tags and appropriate cache policies for edge cacheable fragments. However, for a fragment to be edge cacheable, you must be able to externally request it from the application server. In other words, if a user types the URL in their browser with the appropriate parameters and cookies for the fragment, WebSphere Application Server must return the content for that fragment.

One of the standard J2EE programming architectures is the model-view-controller (MVC) architecture, where a call to a controller servlet might include one or more child JSP files to construct the view. When

using the MVC programming model, the child JSP files are edge cacheable only if you can request these JSP files externally, which is not usually the case. For example, if a child JSP file uses one or more request attributes that are determined and set by the controller servlet, you cannot cache that JSP file on the edge. You can use alternate URL support to overcome this limitation by providing an alternate controller servlet URL used to invoke the JSP file.

The alternate URL for a JSP file or a servlet is set in the `cachespec.xml` file as a property with the name `alternate_url`. You can set the alternate URL either on a per cache-entry basis or on a per cache-id basis. It is valid only if the `EdgeCacheable` property is also set for that entry. If the `EdgeCacheable` property is not set, the `alternate_url` property is ignored. The following is a sample cache policy using the `alternate_url` property:

```
<cache-entry>
 <class>servlet</class>
 <name>/AltUrlTest2.jsp</name>
 <property name="EdgeCacheable">true</property>
 <property name="alternate_url">/alturlcontroller2</property>
 <cache-id>
 <timeout>600</timeout>
 <priority>2</priority>
 </cache-id>
</cache-entry>
```

For more information on the `cachespec.xml` file, see `Cachespec.xml` file.

## Configuring external cache groups

The dynamic cache can control caches outside of the application server, such as IBM Edge Server, an IBM HTTP Server for distributed platforms' Fast Response Cache Accelerator (FRCA) cache, and a WebSphere HTTP Server for distributed platforms plug-in ESI Fragment Processor. When external cache groups are defined, the dynamic cache matches externally cacheable cache entries with those groups, and pushes cache entries and invalidations out to those groups. This allows WebSphere Application Server to manage dynamic content beyond the application server. The content can then be served from the external cache, instead of the application server, improving savings in performance.

1. Open the administrative console.
2. Enable the dynamic cache.
  - a. Click **Servers > Application Servers** in the administrative console navigation tree.
  - b. Click a *server*.
  - c. Click **Dynamic Cache Service**.
  - d. Select the check box in the **Startup state** field to enable the dynamic cache.
3. Define the external cache group in which WebSphere Application Server should control.
  - a. Click **External Caching Groups** from the Dynamic Cache administrative console page.
  - b. Click **New** or choose an external cache group from the list.
4. Configure cache group members.
  - a. Click **External cache groups** from the Dynamic Cache administrative console page. Then click **New** or choose an external cache group from the list.
  - b. Click **External cache group members > New** or choose an external cache group member from the list.
  - c. Type the configuration string in the **Address** field.
  - d. Type the adapter bean name in the **Adapter Bean Name** field.
  - e. **Save** the configuration.
  - f. Click **Apply** or **OK**.



## External cache group collection

Use this page to define sets of external caches controlled by WebSphere Application Server on Web servers, such as IBM Edge Server and IBM HTTP Server.

To view this administrative console page, click **Servers > Application Servers > server > Dynamic Cache Service > External Cache Groups**.

### **Name:**

Specifies the external cache group name.

The external cache group name needs to match the **ExternalCache** property as defined in the servlet or JSP `cachespec.xml` file.

When external caching is enabled, the cache matches pages with its URIs and pushes matching pages to the external cache. The entries can then be served from the external cache, instead of the application server.

### **Type:**

Specifies the external cache group type.

## External cache group settings

Use this page to configure sets of external caches controlled by WebSphere Application Server on Web servers, such as IBM Edge Server and IBM HTTP Server.

To view this administrative console page, click **Servers > Application Server > server > Dynamic Cache Service > External Cache groups > external\_cache\_group**.

### **Name:**

Specifies the external cache group name.

The external cache group name needs to match the **ExternalCache** property as defined in the servlet or JavaServer Pages (JSP) `cachespec.xml` file.

When external caching is enabled, the cache matches pages with its URIs and pushes matching pages to the external cache. The entries can then be served from the external cache, instead of the application server. This ability creates a significant savings in performance.

### **Type:**

Specifies the external cache group type.

## External cache group member collection

Use this page to define specific caches that are members of a cache group.

To view this administrative console page, click **Servers > Application Servers > server > Dynamic Cache Service > External Cache groups > external\_cache\_group > External cache group members**.

### **Address:**

Specifies a configuration string used by external cache adapter bean to connect to the external cache.

### **AdapterBeanName:**

Specifies the adapter bean name.

Example adapter bean names supported in WebSphere Application Server are:

Cache type	AdapterBeanName	Address
Fast Response Cache Accelerator (FRCA) cache - AFPA	com.ibm.ws.cache.servlet.Afpa	Port on which afpa listens
ESI Fragment Processor plug-in for WebSphere HTTP server	com.ibm.websphere.servlet.cache.ESIInvalidatorServlet	local host
IBM Edge Server - WTE	com.ibm.websphere.edge.dynacache.WteAdapter	hostname:port (host name and port on which WTE is listening)

## External cache group member settings

Use this page to configure specific caches that are members of a cache group.

To view this administrative console page, click **Servers > Application Servers > server > Dynamic Cache Service > External Cache groups > external\_cache\_group > External cache group members > external\_cache\_group\_member**.

### **Address:**

Specifies a configuration string used by external cache adapter bean to connect to the external cache.

### **AdapterBeanName:**

Specifies the adapter bean name.

Example adapter bean names supported in WebSphere Application Server are:

<b>AFPA</b>
AdapterBeanName: com.ibm.ws.cache.servlet.Afpa
Address: Port on which afpa listens
<b>ESI</b>
AdapterBeanName: com.ibm.websphere.servlet.cache.ESIInvalidatorServlet
Address: local host
<b>WTE</b>
AdapterBeanName: com.ibm.websphere.edge.dynacache.WteAdapter
Address: hostname:port (host name and port on which WTE is listening)

## Configuring high-speed external caching through the Web server

IBM HTTP Server for Windows NT and Windows 2000 operating systems contains a high-speed cache referred to as the *Fast Response Cache Accelerator*, or *cache accelerator*.

The Fast Response Cache Accelerator is available on Windows NT and Windows 2000 operating systems and AIX platforms. However, support to cache dynamic content is only available on Windows NT and Windows 2000 operating systems.

You can enable cache accelerator to cache static and dynamic content. To enable cache accelerator for caching static content, add the following directives to the `http.conf` configuration file, in the IBM HTTP Server conf directory:

- `AfpaEnable`
- `AfpaCache on`
- `AfpaLogFile "install_root\IBMHttpServer\logs\afpalog" V-ECLF`

To enable cache accelerator for caching dynamic content, such as servlets and Java Server Pages (JSP) files, configure the WebSphere Application Server and the IBM HTTP Server for distributed platforms:

1. Configure WebSphere Application Server to enable Fast Response Cache Accelerator. It is important to follow all the steps for every application server in the cluster.
  - a. Turn on servlet caching for each application server that uses the cache accelerator.
  - b. Configure an external cache group on the application server:
    - 1) Click **Servers > Application Servers > server1 > Dynamic Cache Service > External Cache Groups**.
    - 2) Click **New** on the External cache group administrative console page to define an external cache group named `afpa` for each application server that uses the cache accelerator.
    - 3) In the **External cache group** field, type `afpa` and apply the changes.
  - c. Add a member to the group with an adapter bean name of `com.ibm.ws.cache.servlet.Afpa`.
    - 1) Click **Afpa > External cache group members**.
    - 2) Click **New** on the External cache group members administrative console page.
    - 3) In the **AdapterBean name** field, type `com.ibm.ws.cache.servlet.Afpa`.
    - 4) In the **Address** field, enter an unused port number.
  - d. Add a cache policy in the `cachespec.xml` file for the servlet or JSP file you want to cache. Add the following property to the cache policy:

```
<property name="ExternalCache">afpa</property>
```
2. Enable cache accelerator on the IBM HTTP Server for distributed platforms:
  - a. Add the following directives to the end of the `httpd.conf` file:
    - `AfpaEnable`
    - `AfpaCache on`
    - `AfpaLogFile "install_root\IBMHttpServer\logs\afpalog" V-ECLF`
    - `LoadModule afpaplugin_module install_root/bin/afpaplugin.dll`
    - `AfpaPluginHost WAS_Hostname:port`, where `WAS_Hostname` is the host name of the application server and `port` is the port you specified in the **Address** field while configuring the external cache group member

The `LoadModule` directive loads the IBM HTTP Server plug-in that connects the Fast Response Cache Accelerator to the WebSphere Application Server fragment cache. If multiple IBM HTTP Servers are routing requests to a single application server, add the directives above to the `http.conf` file of each of these IBM HTTP Servers for distributed platforms. If one IBM HTTP Server is routing requests to a cluster of application servers, add the `AfpaPluginHost WAS_Hostname:port` directive to the `http.conf` file for each application server in the cluster. For example, if there are three application servers in the cluster, add the following directives to the `http.conf` file:

- `LoadModule afpaplugin_module install_root/bin/afpaplugin.dll`
- `AfpaPluginHost WAS1_Hostname:port1`
- `AfpaPluginHost WAS2_Hostname:port2`
- `AfpaPluginHost WAS3_Hostname:port3`

### **Configuring Fast Response Cache Accelerator cache size through a distributed platforms Web server:**

In the default IBM HTTP Server for distributed platforms configuration, the maximum Fast Cache Accelerator dynamic cache size is calculated as 1/8 of physical pin-able memory. On a machine with 384

megabytes of RAM, it allows a maximum of approximately 50 megabytes for the Fast Cache Accelerator dynamic cache. When this limit is reached, the Cache Accelerator then deletes older entries to cache new ones.

Follow these steps to configure the Cache Accelerator:

Using the IBM HTTP Server for distributed platforms' `AfpaDynaCacheMax` directive, tune the maximum allowed cache size:

1. Place the directive in the global server configuration scope, along with the other default Fast Cache Accelerator directives.
2. Enable Fast Cache Accelerator. To enable the Fast Cache Accelerator, update the following directives in this IBM HTTP Server's `http.conf` file:

```
AfpaEnable
AfpaCache on
AfpaLogFile "c:/Program Files/IBM HTTP Server/logs/afpalog" V-ECLF
AfpaDynaCacheMax 10
```

These above settings limit the dynamic cache size to 10 megabytes. If you use these directives to increase cache size, do not make the cache so large that all the physical memory is consumed. Determine how much memory is available when all applications are running, by using the Windows Task Manager.

Assign no more than 50% of available physical memory to the dynamic cache. Specifying too large a cache not only decreases the performance of other applications, but also puts you at a risk for completely running out of memory.

The default configuration does not include the `AfpaDynaCacheMax` directive where the cache size is automatically calculated as 1/8 of physical memory.

---

## Displaying cache information

The dynamic cache monitor is an installable Web application that displays simple cache statistics, cache entries, and cache policy information.

1. Use the administrative console to install the cache monitor application from the `install_root/installableApps` directory. The application is named `CacheMonitor.ear`. Install the cache monitor onto the application server you are trying to monitor. Installing the cache monitor on the `admin_host` (port 909x) is more secure than installing it on the `default_host` (908x), and so it is preferable to install it onto the `admin_host`.
2. Configure the HTTP transport and host alias for the server with cache monitor installed.
  - a. If you installed the cache monitor on the `admin_host` (port 909x), check if an HTTP transport has been created. Click **Application Servers > server > Web Container > HTTP Transports**. If an HTTP transport (port 909x) does not exist you must create an HTTP transport in the `admin_host` for this server. If you are using `server1`, an HTTP transport is installed by default for admin port 9090.
  - b. Add a host alias for the port your server is using. Click **Environment > Virtual Hosts > host\_type > Host Aliases** and create a new **Host Name** and **Port** to add to the list.
  - c. You can then access the cache monitor using `http://your_host_name:your_port_number/cachemonitor`.

**Tip:** You can find the port number in the `SystemOut.log` file. Look for message SRVE0171I.

3. Access the Web application using a Web browser and the URL `http://your_host_name:your_port_number/cachemonitor`, where `your_port_number` is the port associated with the host on which you installed the cache monitor application.
4. Verify that the cache monitor is working properly.
  - a. View the Statistics page and verify the cache configuration and cache data. Click **Reset Statistics** to reset the counters

- b. View the Cache Policies page to see which cache policies are currently loaded in the dynamic cache. Click on a template to view the cache ID rules for the template.
- c. View the Cache Contents page to examine the contents currently cached in memory.
- d. View the ESI Statistics page to view data about the current ESI processors configured for caching. Click **Refresh Statistics** to see the latest statistics or content from the ESI processors. Click **Reset Statistics** to reset the counters.
- e. View the Disk Offload page to view content currently off-loaded from memory to disk.

When viewing contents on memory or disk, click on a template to view all entries for that template, click on a dependency ID to view all entries for the ID, or click on the cache ID to view the entire data cached for that entry.

5. Use the cache monitor to perform basic operations on data in a cache.

**Remove an entry from cache**

Click **Invalidate** when viewing a cache entry.

**Remove all entries for a certain dependency ID**

Click **Invalidate** when viewing entries for a dependency ID.

**Remove all entries for a certain template**

Click **Invalidate** when viewing entries for a template.

**Move an entry to the front of the Least Recently Used queue to avoid eviction**

Click **Refresh** when viewing a cache entry.

**Move an entry from disk to cache**

Click **Send to Memory** when viewing a cache entry on disk.

**Clear the entire contents of the cache**

Click **Clear Cache** while viewing statistics or contents.

**Clear the contents on the ESI processors**

Click **Clear Cache** while viewing ESI statistics or contents.

**Clear the contents of the disk cache**

Click **Clear Disk** while viewing disk contents.

---

## Configuring cacheable objects with the cachespec.xml file

Define cacheable objects inside the `cachespec.xml`, found inside the Web module `WEB-INF` or enterprise bean `META-INF` directory.

You can place a global `cachespec.xml` in the application server properties directory, but the recommended method is to place the cache configuration file with the deployment module. The root element of the `cachespec.xml` file is `<cache>`, which contains `<cache-entry>` elements.

Within a `<cache-entry>...</cache-entry>` element are parameters that allow you to complete the following tasks to enable the dynamic cache with the `cachespec.xml` file:

1. Develop a `cachespec.xml` file.
  - a. Create a caching configuration file.
 

In the `<install_root>/properties` directory, locate the `cachespec.sample.xml` file.
  - b. Copy the `cachespec.sample.xml` file to `cachespec.xml` in Web module `WEB-INF` or enterprise bean `META-INF` directory.
2. Define the cache-entry elements necessary to identify the cacheable objects. See the topic “Cachespec.xml file” on page 1356 for a list of elements.
3. Develop cache-ID rules.

To cache an object, WebSphere Application Server must know how to generate unique IDs for different invocations of that object. The `<cache-id>` element performs that task. Each cache entry can have multiple cache-ID rules that execute in order until either a rule returns non-empty cache-ID or no more rules remain to execute. If none of the cache-ID generation rules produce a valid cache ID, then the object is not cached. Develop these IDs in one of two ways:

- Use the `<component>` element defined in the cache policy of a cache entry (recommended)
- Write custom Java code to build the ID from input variables and system state

To configure the cache entry to use the `IdGenerator`, specify your `IdGenerator` in the XML file by using the `<idgenerator>` tag, for example:

```
<cache-entry>
 <class>servlet</class>
 <name>/servlet/CommandProcessor</name>
<cache-id>
 <idgenerator>com.mycompany.SampleIdGeneratorImpl</idgenerator>
 <timeout>60</timeout>
</cache-id>
</cache-entry>
```

4. Specify dependency ID rules. Use dependency ID elements to specify additional cache group identifiers that associate multiple cache entries to the same group identifier.

The dependency ID is generated by concatenating the dependency ID base string with the values returned by its component elements. If a required component returns a null value, then the entire dependency ID does not generate and is not used. You can validate the dependency IDs explicitly through the WebSphere Dynamic Cache API, or use another cache-entry `<invalidation>` element. Multiple dependency ID rules can exist per cache-entry. All dependency ID rules separately execute. See the topic “Cachespec.xml file” on page 1356 for a list of `<component>` elements.

5. Invalidate other cache entries as a side effect of this object execution, if relevant. You can define invalidation rules in exactly the same manner as dependency IDs. However, the IDs that generate by invalidation rules are used to invalidate cache entries that have those same dependency IDs.

The invalidation ID is generated by concatenating the invalidation ID base string with the values returned by its component element. If a required component returns a null value, then the entire invalidation ID is not generated and no invalidation occurs. Multiple invalidation rules can exist per cache-entry. All invalidation rules separately execute.

6. Verify the cacheable page.

Typically you declare several `<cache-entry>...</cache-entry>` elements inside a `cachespec.xml` file.

The dynamic cache responds to changes in this file. When new versions of the `cachespec.xml` are detected, the old policies are replaced. Objects cached through the old policy file are not automatically invalidated from the cache; they are either reused with the new policy or eliminated from the cache through its replacement algorithm.

For each of the three IDs (cache, dependency, invalidation) generated by cache entries, a `<cache-entry>` can contain multiple elements. The dynamic cache will execute the `<cache-id>` rules in order, and the first one that successfully generates an ID will be used to cache that output. If the object is to be cached, each one of the `<dependency-id>` elements will be executed to build a set of dependency IDs for that cache entry. Finally, each of the `<invalidation>` elements will be executed, building a list of IDs that the dynamic cache will invalidate, whether or not this object is cached.

## Verifying the cacheable page

Verify the cacheable page by following these steps:

1. View the snoop servlet in the default application by accessing the URI: `/snoop`
2. Invoke and reload the URI several times using a different Web browser or using different parameters. This action returns the same output for the snoop servlet. The snoop servlet is now operating incorrectly, because it displays the request information from its first invocation rather than from the current request.
3. Inspect the entry in the cache with the dynamic cache monitor.



## Cachespec.xml file

The cache parses the `cachespec.xml` file on startup, and extracts from each `<cache-entry>` element a set of configuration parameters. Every time a new servlet or other cacheable object initializes, the cache attempts to match each of the different cache-entry elements, to find the configuration information for that object. Different cacheable objects have different `<class>` elements. You can define the specific object a cache policy refers to using the `<name>` element.

### Location

The `cachespec.xml` file is found inside the `WEB-INF` directory of a Web module.

You can place a global `cachespec.xml` file in the application server properties directory, but the recommended method is to place the cache configuration file with the deployment module. (To place the cache configuration file with the deployment module, use the Assembly ToolkitApplication Assembly Tool (AAT) to define the cacheable objects.

The root element of the `cachespec.xml` file is `cache`, which contains `cache-entry` elements.

The `cachespec.dtd` file is available in the application server properties directory.

### Usage notes

#### Cachespec.xml elements

Each cache entry must specify certain basic information that the dynamic cache uses to process that entry. This section explains the function of each cache entry element of the `cachespec.xml` file including:

- `class`
- `name`
- `sharing-policy`
- `property`
- `cache-id`

#### class

```
<class>command | servlet | webservice</class>
```

This element is required and governs how the application server interprets the remaining cache policy definition. The value `servlet` refers to servlets and JavaServer Pages (JSP) files deployed in the WebSphere Application Server servlet engine. The `webservice` class extends the `servlet` with special component types for Web services requests. Finally, the value `command` refers to classes using the WebSphere command programming model. The following examples illustrate the `class` element:

```
<class>command</class>
<class>servlet</class>
<class>webservice</class>
```

#### name

```
<name>name</name>
```

where *name* is the fully qualified class name of the command, servlet, or Web service.

There are two ways to use `<name>` to specify a cacheable object:

- For commands and Web services, this required element must include the package name, if any, and class name, including a trailing `.class`, of the configured object.
- For Web Services, include the URI of the SOAP router associated with the Web service that you want to cache.



- For servlets and JSP files, if the cachespec.xml file is in the WebSphere Application Server properties directory, this required element must include the full URI of the JSP file or servlet to cache. For servlets and JSP files, if the cachespec.xml file is in the Web application, this required element can be relative to the specific Web application context root.

**Note:** The preferred location of the cachespec.xml file is in the Web application, not the properties directory.

You can specify multiple <name> elements within a <cache-entry> if you have different mappings that refer to the same servlet.

The following examples illustrate the name element:

```
<name>com.mycompany.MyCommand.class</name>
<name>default_host:/servlet/snoop</name>
<name>com.mycompany.beans.MyJavaBean</name>
<name>mywebapp/myjsp.jsp</name>
```

### sharing-policy

```
<sharing-policy> not-shared | shared-push | shared-pull | shared-push-pull</sharing-policy>
```

When working within a cluster with a distributed cache, these values determine the sharing characteristics of entries created from this object. If this element is not present, a not-shared value is assumed. Also, in non-distributed environments, not-shared is the only valid value. When enabling replication, the default value is shared-push only. This property does not affect distribution to Edge Side Include processors through the Edge fragment caching property.

Value	Description
not-shared	Cache entries for this object are not shared among different application servers. These entries can contain non-serializable data. For example, a cached servlet can place non-serializable objects into the request attributes, if the <class> type supports it.
shared-push	Cache entries for this object are automatically distributed to the dynamic caches in other application servers or cooperating Java virtual machines (JVMs). Each cache has a copy of the entry at the time it is created. These entries cannot store non-serializable data.
shared-pull	Cache entries for this object are shared between application servers on demand. If an application server gets a cache miss for this object, it queries the cooperating application servers to see if they have the object. If no application server has a cached copy of the object, the original application server executes the request and generates the object. These entries cannot store non-serializable data. This mode of sharing is not recommended.
shared-push-pull	Cache entries for this object are shared between application servers on demand. When an application server generates a cache entry, it broadcasts the cache ID of the created entry to all cooperating application servers. Each server then knows whether an entry exists for any given cache ID. On a given request for that entry, the application server knows whether to generate the entry or pull it from somewhere else. These entries cannot store non-serializable data.

The following example shows a sharing policy:

```
<sharing-policy>not-shared</sharing-policy>
```

### property

```
<property name="key">value</property>
```

where *key* is the name of the property defined for this cache entry element, and *value* is the corresponding value.

You can set optional properties on a cacheable object, such as a description of the configured servlet. The class determines valid properties of the cache entry. At this time, the following properties are defined:

Property	Valid classes	Value
ApplicationName	All	Overrides the J2EENAME application ID so that multiple applications can share a common cache ID namespace.
EdgeCacheable	Servlet	True or false. Default is false. If the property is true, then the given servlet or JSP file is externally requested from an Edge Side Include processor. Whether or not the servlet or JSP file is cacheable depends on the rest of the cache specification.
ExternalCache	Servlet	Specifies the external cache name. The external cache name needs to match the external cache group name.
consume-subfragments	Servlet or Web service	True or false. Default is false. When a servlet is cached, only the content of that servlet is stored, and includes placeholders for any other fragments to which it includes or forwards. Consume-subfragments (CSF) tells the cache not to stop saving content when it includes a child servlet. The parent entry, the one marked CSF, includes all the content from all fragments in its cache entry, resulting in one big cache entry that has no includes or forwards, but the content from the whole tree of entries. This can save a significant amount of application server processing, but is typically only useful when the external HTTP request contains all the information needed to determine the entire tree of included fragments.
alternate_url	Servlet	Specifies the alternate URL used to invoke the servlet or JSP file. The property is valid only if the EdgeCacheable property also is set for the cache entry.

persist-to-disk	All	True or false. Default is true. When this property is set to false, the cache entry is not written to the disk when overflow or server stopping occurs.
save-attributes	Servlet	True or false. Default is true. When this property is set to false, the request attributes are not saved with the cache entry.

## cache-id

To cache an object, the application server must know how to generate a unique ID for different invocations of that object. These IDs are built either from user-written custom Java code or from rules defined in the cache policy of each cache entry. Each cache entry can have multiple cache ID rules that are executed in order until either:

- A rule returns a non-empty cache ID, or
- No more rules are left to execute.

If none of the cache ID generation rules produce a valid cache ID, the object is not cached.

Each `cache-id` element defines a rule for caching an object and is composed of the sub-elements `component`, `timeout`, `priority`, and `property`. The following example illustrates a `cache-id`:

```
<cache-id>component*| timeout? | priority? | property* </cache-id>
```

## component sub-element

Use the `component` sub-element to generate a portion of the cache ID. Each `component` sub-element consists of the attributes `id`, `type`, and `ignore-value`, and the elements `method`, `field`, `required`, `value`, and `not-value`.

- Use the `id` attribute to identify the component.
- Use the `type` attribute to identify the type of component. The following table lists the values for the `type`.

Type	Valid classes	Meaning
method	command	Calls the indicated method on the command or object
field	command	Retrieves the named field in the command or object
parameter	servlet	Retrieves the named parameter value from the request object
parameter-list	servlet	Retrieves a list of values for the named parameter
session	servlet	Retrieves the named value from the <code>HTTPSession</code>
cookie	servlet	Retrieves the named cookie value
attribute	servlet	Retrieves the named request attribute
header	servlet and Web service	Retrieves the named request header
pathInfo	servlet	Retrieves the <code>pathInfo</code> from the request
servletpath	servlet	Retrieves the servlet path
locale	servlet	Retrieves the request locale

SOAPEnvelope	Web service	Retrieves the SOAPEnvelope from a Web services request. An ID attribute of Hash uses a Hash of the SOAPEnvelope, while Literal uses the SOAPEnvelope as received.
SOAPAction	Web service	Retrieves the SOAPAction header, (if available), for a Web services request.
serviceOperation	Web service	Retrieves the service operation for a Web services request
serviceOperationParameter	Web service	Retrieves the specified parameter from a Web services request

- Use the `ignore-value` attribute to specify whether or not to use the value returned by this component in cache ID formation. This is an optional attribute with a default value of `false`. If the value is `true`, only the ID of the component is used when creating a cache ID, or no output is used when creating a dependency or invalidation ID.

- Use the **method** element to call a void method on a returned object. You can infinitely nest method and field objects in any combination. The method must be public and is not valid for edge-cacheable components. For example:

```
<component id="getUser" type="method"><method>getUserInfo
<method>getName</method></method></component>
```

This method is equivalent to `getUser().getUserInfo().getName()`

- Use the **field** element to access a field in a returned object. You can infinitely nest method and field objects in any combination. The field must be public. Not valid for edge-cacheable components. For example:

```
<component id="getUser" type="method"><method>getUserInfo
<field>name</field></method></component>
```

This method is equivalent to `getUser().getUserInfo().name`

- Use the **required** element to specify whether or not this component must return a non-null value for this cache ID for it to represent a valid cache. If set to `true`, this component must return a non-null value for this cache ID to represent a valid cache ID. If set to `false`, the default, a non-null value is used in the formation of the cache ID and a null value means that this component is not used at all in the ID formation. For example:

```
<required>true</required>
```

- Use the **value** element to specify values that must match to use this component in cache ID formation. For example:

```
<component id="getUser" type="method"><value>blue</value>
<value>red</value> </component>
```

- Use the **not-value** element to specify values that must not match to use this component in cache ID formation. This method is similar to `<value>`, but instead prescribes the defined values from caching. You can use multiple `<not-value>` elements when there is more than one invalid value. For example:

```
<component id="getUser" type="method">
<required>true</required>
<not-value>blue</not-value>
<not-value>red</not-value></component>
```

The component element can have either a method and a field element, a value element, or a not-value element. The method and field elements apply only to commands. The following example illustrates the attributes of a component element:

```
<component id="isValid" type="method" ignore-value="true"><component>
```

### timeout sub-element

The timeout sub-element is used to specify a time-to-live (TTL) value for the cache entry. For example,

`<timeout>value</timeout>`

where *value* is the amount of time, in seconds, to keep the cache entry. If 0, or a negative value is specified, the cache entry is kept indefinitely.

### priority sub-element

Use the priority sub-element to specify the priority of a cache entry in a cache. The priority weighting is used by the least recently used (LRU) algorithm, of the cache to decide which entries to remove from the cache if the cache runs out of storage space. For example,

`<priority>value</priority>`

where *value* is a positive integer between 1 and 255 inclusive.

### property sub-element

Use the property sub-element to specify generic properties for the cache entry. For example,

`<property name="key">value</property>`

where *key* is the name of the property to define, and *value* is the corresponding value.

For example:

`<property name="description">The Snoop Servlet</property>`

Property	Valid classes	Meaning
sharing-policy/timeout/priority	All	Overrides the settings for the containing cache entry when the request matches this cache ID.
EdgeCacheable	servlet	Overrides the settings for the containing cache entry when the request matches this cache ID.

### idgenerator and metadatagenerator elements

Use the `idgenerator` element to specify the class name loaded for the generation of the cache ID. The `IdGenerator` must implement the `com.ibm.websphere.servlet.cache.IdGenerator` interface. The `IdGenerator` must build and set cache IDs, group IDs and invalidation IDs. An example of the `idgenerator` element follows:

`<idgenerator> classname classname </idgenerator>`

(where *classname*= Fully qualified name of the class to use)

Use the `metadatagenerator` element to specify the class name loaded for the metadata generation cache ID. The `MetadataGenerator` class must implement the `com.ibm.websphere.servlet.cache.MetadataGenerator` interface. The `MetadataGenerator` defines properties like timeout, external caching or generic properties. An example of the `metadatagenerator` element follows:

`<metadatagenerator> classname classname </metadatagenerator>`

(where *classname*= Fully qualified name of the class to use)

---

## Configuring command caching

Cacheable commands are stored in the cache for re-use with a similar mechanism for servlets and Java Server Pages (JSP) files. However, in this case, the unique cache IDs are generated based on methods and fields present in the command as input parameters. For example, a **GetStockQuote** command can have a symbol as its input parameter.

A unique cache ID can generate from the name of the command, plus the value of the symbol.

To use command caching you must:  
Create a command.

1. Define an interface. The Command interface specifies the most basic aspects of a command. You must define the interface that extends one or more of the interfaces in the command package. The command package consists of three interfaces:
  - TargetableCommand
  - CompensableCommand
  - CacheableCommand

In practice, most commands implement the TargetableCommand interface, which allows the command to execute remotely. The code structure of a command interface for a targetable command follows:

```
...
import com.ibm.websphere.command.*;
public interface MyCommand extends TargetableCommand {
 // Declare application methods here
}
```

2. Provide an implementation class for the interface. Write an interface that extends the CacheableCommandImpl class and implements your command interface. This class contains the code for the methods in your interface, the methods inherited from extended interfaces like the CacheableCommand interface, and the required or abstract methods in the CacheableCommandImpl class. You can also override the default implementations of other methods provided in the CacheableCommandImpl class.

## Command class

To write a command interface, extend one or more of the three interfaces included in the command package. The base interface for all commands is the Command interface. This interface provides only the client-side interface for generic commands and declares three basic methods:

- **isReadyToCallExecute**. This method is called on the client side before the command passes to the server for execution.
- **execute**. This method passes the command to the target and returns any data.
- **reset**. This method reverts any output properties to the values they had before the execute method was called so that you can reuse the object.

The implementation class for your interface must contain implementations for the isReadyToCallExecute and reset methods.

## CacheableCommandImpl class

Commands are implemented by extending the class CacheableCommandImpl, which implements the CacheableCommand interface.

The CacheableCommandImpl class is an abstract class that provides implementations for some of the methods in the CacheableCommand interface, for example, setting return values. This class declares additional methods that the application must implement, for example, how to execute the command.

The code structure of an implementation class for the CacheableCommand interface follows:

```
...
import com.ibm.websphere.command.*;
public class MyCommandImpl extends CacheableCommandImpl
implements MyCommand {
 // Set instance variables here ...
 // Implement methods in the MyCommand interface ...
 // Implement abstract methods in the CacheableCommandImpl class
 ...
}
```

## Example: Caching a command object

This example of command caching is a simple stock quote command.

The following is a stock quote command bean. It accepts a ticker as an input parameter and produces a price as its output parameter.

```
public class QuoteCommand extends CacheableCommandImpl
{
 private String ticker;
 private double price;
 // called to validate that command input parameters have been set
 public boolean isReadyToCallExecute() {
 return (ticker!=null);
 }
 // called by a cache-hit to copy output properties to this object
 public void setOutputProperties(TargetableCommand fromCommand) {
 QuoteCommand f = (QuoteCommand)fromCommand;
 this.price = f.price;
 }

 // business logic method called when the stock price must be retrieved
 public void performExecute()throws Exception {...}

 //input parameters for the command
 public void setTicker(String ticker) { this.ticker=ticker;}
 public String getTicker() { return ticker;}

 //output parameters for the command
 public double getPrice() { return price;};
}
```

To cache the above command object using the stock ticker as the cache key and using a 60 second time-to-live, use the following cache policy:

```
<cache>
<cache-entry>
<class>command</class>
<sharing-policy>not-shared</sharing-policy>
<name>QuoteCommand</name>
<cache-id>
<component type="method" id="getTicker">
<required>true</required>
</component>
<priority>3</priority>
<timeout>60</timeout>
</cache-id>
</cache-entry>
</cache>
```



---

## Using the DistributedMap interface for the dynamic cache

The DistributedMap interface is a simple interface for the dynamic cache. Using the DistributedMap interface, J2EE applications and system components can cache and share Java objects by storing a reference to the object in the cache. The default WebSphere dynamic cache instance is created if dynamic cache service is enabled in the administrative console. This default instance is bound into the global JNDI namespace using the name `services/cache/distributedmap`. For information on how to enable the dynamic cache service globally, see [Enabling the dynamic cache service](#).

Multiple instances of the DistributedMap interface on the same Java virtual machine (JVM) enables applications to separately configure cache instances as needed. Each instance of DistributedMap has its own properties:

- **Cache JNDI name.** The JNDI name is used for calling a particular DistributedMap instance.
- **Cache size (Property name: `cacheSize`, Default: 2000).** Dynamic cache maintains the maximum number of entries in memory.
- **Enable disk offload (Property name: `enableDiskOffload`, Default: `false`).**
  - Disk offload is disabled - If a new entry is created while the cache is full, the priorities configured for each entry, along with at least recently used algorithm are used to remove the entry from the cache in memory.
  - Disk offload is enabled - If a new entry is created while the cache is full, the priorities configured for each entry, along with a recently used algorithm are used to remove the entry from the cache in memory. The entry that is removed is copied to the local file system, specified by the disk offload location.
- **Disk offload location (Property name: `diskOffloadLocation`).** If disk offload location is not specified, the default location, `$install_root/temp/node/servername/_dynacache/cacheJNDIname` will be used. If disk offload location is specified, the node, server name, and cache instance name will be appended. For example, `$install_root/diskoffload` generates the location as `$install_root/diskoffload/node/servername/cacheJNDIname`. This value is ignored if `enableDiskOffload` is `false`.
- **Flush to disk default (Property name: `flushToDiskOnStop`, Default: `false`).** Set this value to `true` to save in-memory cached objects to disk when the server is stopped. This value is ignored if `enableDiskOffload` is `false`.
- **Use listener context (Property name: `useListenerContext`, Default: `false`).** Used for cache invalidation listener. If the user wants to use listener J2EE context for callback, set this to `true`. If the user wants to use the caller thread context for callback, set this to `false`.

There are three methods for configuring and using cache instances.

### Method 1 - Administrative Console

You can create additional cache instances using the Administrative Console. Click **Resources > Cache Instances**. After you create a new cache instance, you can optionally add additional properties, such as `flushToDiskOnStop` and `useListenerContext`, using the custom properties. See “Cache instance settings” on page 1366 for more information.

If you defined two cache instances in the Administrative Console with JNDI names of **`services/cache/instance_one`** and **`services/cache/instance_two`**, you can use the following code to look up the cache instances:

```
InitialContext ic = new InitialContext();
DistributedMap dm1 = (DistributedMap)ic.lookup("services/cache/instance_one");
DistributedMap dm2 = (DistributedMap)ic.lookup("services/cache/instance_two");
```

**Note:** For more information about the DistributedMap interface, see the Javadoc for the `com.ibm.websphere.cache` package.

## Method 2 - Properties file

**Note:** The `distributedmap.properties` file will be deprecated in a future release. Using the Administrative Console is the recommended way to configure cache instances.

You can create additional cache instances using the `distributedmap.properties` file with the following format :

```
cache.instance.0=/services/cache/instance_one
cache.instance.0.cacheSize=1000
cache.instance.0.enableDiskOffload=true
cache.instance.0.diskOffloadLocation=${WAS_INSTALL_ROOT}/diskOffload
cache.instance.0.flushToDiskOnStop=true
cache.instance.0.useListenerContext=true
cache.instance.1=/services/cache/instance_two
cache.instance.1.cacheSize=1500
cache.instance.1.enableDiskOffload=false
cache.instance.1.flushToDiskOnStop=false
cache.instance.1.useListenerContext=false
```

The preceding example creates two `DistributedMap` instances named `instance_one` and `instance_two`. `instance_one` has a cache entry size of 1,000 and `instance_two` has a cache entry size of 1,500. Disk offload is enabled in `instance_one` and disabled in `instance_two`. Use listener context is enabled in `instance_one` and disabled in `instance_two`. Flush to disk on stop is enabled in `instance_one` and disabled in `instance_two`.

You must place the `distributedmap.properties` file in either your application server or application class path. For example, you can use your application WAR file, `WEB-INF\classes` directory, or `was_root\classes` directory. The first entry in the properties file (`cache.instance.0`) specifies the JNDI name for the cache instance in the global namespace. You can use the following code to look up the cache instance:

```
InitialContext ic = new InitialContext();
DistributedMap dm1 = (DistributedMap)ic.lookup("services/cache/instance_one");
DistributedMap dm2 = (DistributedMap)ic.lookup("services/cache/instance_two");
```

**Note:** For more information about the `DistributedMap` interface, see the Javadoc for the `com.ibm.websphere.cache` package.

## Method 3 - Resource references

**Note:** Method three is an extension to method one or method two, listed above. First use either method one or method two.

Define a `resource-ref` in your module deployment descriptor (`web.xml` and `ibm-web-bnd.xml` files) and look up the cache using the `java:comp` namespace.

### Resource-ref example:

#### File: `web.xml`

```
<resource-ref id="ResourceRef_1">
 <res-ref-name>dmap/LayoutCache</res-ref-name>
 <res-type>com.ibm.websphere.cache.DistributedMap</res-type>
 <res-auth>Container</res-auth>
 <res-sharing-scope>Shareable</res-sharing-scope>
</resource-ref>
<resource-ref id="ResourceRef_2">
 <res-ref-name>dmap/UserCache</res-ref-name>
 <res-type>com.ibm.websphere.cache.DistributedMap</res-type>
 <res-auth>Container</res-auth>
 <res-sharing-scope>Shareable</res-sharing-scope>
</resource-ref>
```

#### File: `ibm-web-bnd.xml`

```
<?xml version="1.0" encoding="UTF-8"?>
```

```

<webappbnd:WebAppBinding xmi:version="2.0" xmlns:xmi="http://www.omg.org/XMI"
xmlns:webappbnd="webappbnd.xmi"
xmlns:webapplication="webapplication.xmi" xmlns:commonbnd="commonbnd.xmi"
xmlns:common="common.xmi"
xmi:id="WebApp_ID_Bnd" virtualHostName="default_host">
 <webapp href="WEB-INF/web.xml#WebApp_ID"/>
 <resRefBindings xmi:id="ResourceRefBinding_1" jndiName="services/cache/instance_one">
 <bindingResourceRef href="WEB-INF/web.xml#ResourceRef_1"/>
 </resRefBindings>
 <resRefBindings xmi:id="ResourceRefBinding_2" jndiName="services/cache/instance_two">
 <bindingResourceRef href="WEB-INF/web.xml#ResourceRef_2"/>
 </resRefBindings>
</webappbnd:WebAppBinding>

```

The following example shows how to look up the resource-ref:

```

InitialContext ic = new InitialContext();
DistributedMap dm1a =(DistributedMap)ic.lookup("java:comp/env/dmap/LayoutCache");
DistributedMap dm2a =(DistributedMap)ic.lookup("java:comp/env/dmap/UserCache");

```

The previous resource-ref example maps `java:comp/env/dmap/LayoutCache` to `/services/cache/instance_one` and `java:comp/env/dmap/UserCache` to `/services/cache/instance_two`. In the examples, `DistributedMap dm1` and `dm1a` are the same map. `DistributedMap dm2` and `dm2a` are the same map.

**Note:** The `DistributedMap` interface does not currently have authorization or access control associated with any of the cache entries.

To learn how to share cached objects in a clustered environment, see [Sharing cached objects in a clustered environment](#).

## Sharing cached objects in a clustered environment

In a clustered environment, the content you place in cache might be shared with other servers in the cluster. The content might also be off-loaded to disk. If you intend to have the cached objects shared or off-loaded to disk, you must make these particular objects serializable. If the objects you place in cache are non-serializable, you must specify that the sharing policy for these objects is "not shared". The `DistributedMap` interface Javadoc contains information about how to specify the sharing policy for a cached object. Specifying a sharing policy other than "not shared" for non-serializable objects can result in poor system performance.

## Cache instance settings

A cache instance is a location, in addition to the default shared dynamic cache, where any Java 2 platform, Enterprise Edition (J2EE) application can store, distribute, and share data. The cache instance gives applications greater flexibility and better tuning of the cache resources. The programming interface to access this cache instance is called `DistributedMap`. See the `DistributedMap` API documentation in the WebSphere Application Server Javadoc for more information.

To view this administrative console page, click **Resources > Cache Instances**.

### Name

Specifies the required display name for the resource.

### JNDI name

Specifies the Java Naming and Directory Interface (JNDI) name for the resource. Use this name when looking up a reference to this cache instance. The returned object has a `DistributedMap` type.

### Description

Specifies a description for the resource. This field is optional.

## Category

Specifies a category string to classify or group the resource. This field is optional.

## Cache Size

Specifies a positive integer for the maximum number of entries that the cache holds.

The range is 100 to 200,000.

## Default Priority

Specifies the default priority for servlets that can be cached. This value determines how long an entry stays in a full cache.

The recommended value is one.

## Enable Disk Offload

Selecting this option will enable disk offloading.

## Disk Offload Location

Specifies the directory used for disk offload.

## Flush to Disk

Specifies if in-memory cached objects are saved to disk when the server is stopped. This value is ignored if the Enable Disk Offload value is false.

## Use Listener Context

Set this value to `true` to have invalidation events sent to registered invalidation listeners using the J2EE context of the listener.

## Additional custom properties (available only when a cache instance is created)

Custom properties are available for the `flushToDiskOnStop` and `useListenerContext` instances.

You can define custom properties after you create a cache instance.

**Name** Specifies the name of the custom property. In this case, the name is either `flushToDiskOnStop` or `useListenerContext`. This field is required.

**Value** Specify `true` or `false` to indicate the value of the custom property. This field is required.

### Description

Describes any bounds or values of the custom property. This field is optional.

## Cache instance collection

Use this page to configure and manage cache instances.

To view this administrative console page, click **Resources > Cache Instances > New**.

## Scope

Use the Scope to limit the availability of the cache instance to a particular cell, node, or server. Select the Cell scope to view and configure cache instances that are available to all servers within the cell. Specify the Node scope to view and configure cache instances that are available to all servers with the particular node. Specify the Server scope to view and configure cache instances that are available only on the specific server.

## Name

Specifies the required display name for the resource.

## JNDI name

Specifies the Java Naming and Directory Interface (JNDI) name for the resource. Use this name when looking up a reference to this cache instance. The returned object is of the DistributedMap type.

## Description

Specifies a description for the resource. This field is optional.

## Category

Specifies a category string to classify or group the resource. This field is optional.

## Cache Size

Specifies a positive integer for the maximum number of entries that the cache holds.

The range is 100 to 200,000.

## Enable Disk Offload

Selecting this option enables disk offloading.

By default, dynamic cache only maintains the number of entries configured in memory. If new entries are created while the cache is full, the priorities configured for each cache entry, along with a least recently used (LRU) algorithm, are used to remove entries from the cache. As an alternative, you can configure disk offload to copy entries onto the file system for access at a later time. The location to copy the entries is configurable.

## Cache instance service settings

The Cache Instance Service enables server-level configuration for Cache Instance resources.

To view this administrative console page, click **Servers > Application Servers > server1 > Cache Instance Service**.

## Startup

Specifies whether the server will attempt to start the specified service when the server starts. Checking the box turns this service on or off.

## Invalidation listeners

Invalidation listener mechanism uses Java events for alerting applications when contents are removed from the cache.

Applications implement the InvalidationListener interface (defined in the `com.ibm.websphere.cache` package) and register it to the cache using the DistributedMap interface. Listeners receive InvalidationEvents (defined in the `com.ibm.websphere.cache` package) when entries from the cache are evicted, due to an explicit user invalidation, timeout, least recently used (LRU) eviction, cache clear, or disk timeout. Applications can immediately recalculate the invalidated data and prime the cache before the next user request.

Enable listener support in DistributedMap before registering listeners. DistributedMap can also be configured to use the invalidation listener J2EE context from registration time during callbacks. Setting the value of the custom property `useListenerContext` to true will enable the invalidation listener J2EE context for callbacks. See Cache instance settings for more information.

The following example shows how to set up an invalidation listener:

```
dmap.enableListener(true); // Enable cache invalidation listener.
InvalidationListener listener = new MyListenerImpl(); //Create invalidation listener object.
dmap.addInvalidationListener(listener); //Add invalidation listener.
:
:
```

```
 :
dmap.removeInvalidationListener(listener); //Remove the invalidation listener.
//This increases performance.
dmap.enableListener(false); // Disable cache invalidation listener.
//This increases performance.
```

For more information about invalidation listeners, see the Javadoc for the `com.ibm.websphere.cache` package.

---

## Example: Caching Web services

The following is an example of building a set of cache policies for a simple Web services application. The application in this example stores stock quotes, and has operations to read, update the price of, and buy a given stock symbol.

Following are two SOAP message examples that the application can receive, with accompanying HTTP Request headers.

The first message sample contains a SOAP message for a `GetQuote` operation, requesting a quote for IBM. This is a read-only operation that gets its data from the back-end, and is very cacheable. In this example the SOAP message is cached and a timeout is placed on its entries to guarantee the quotes it returns are not too out of date.

### Message example 1

```
POST /soap/servlet/soaprouter
HTTP/1.1
Host: www.myhost.com
Content-Type: text/xml; charset="utf-8"
SOAPAction: urn:stockquote-lookup
<SOAP-ENV:Envelope xmlns:SOAP-ENV="http://schemas.xmlsoap.org/soap/envelope/"
SOAP-ENV:encodingStyle="http://schemas.xmlsoap.org/soap/encoding/">
<SOAP-ENV:Body>
<m:getQuote xmlns:m="urn:stockquote:>
<symbol>IBM</symbol>
</m:getQuote>
</SOAP-ENV:Body>
</SOAP-ENV:Envelope>
```


The `SOAPAction` HTTP header in the request is defined in the SOAP specification and is used by HTTP proxy servers to dispatch requests to particular HTTP servers. WebSphere Application Server dynamic cache can use this header in its cache policies to build IDs without having to parse the SOAP message.

Message example 2 illustrates a SOAP message for a `BuyQuote` operation. While message 1 is cacheable, this message is not, because it updates the back-end database.

### Message example 2

```
POST /soap/servlet/soaprouter
HTTP/1.1
Host: www.myhost.com
Content-Type: text/xml; charset="utf-8"
SOAPAction: urn:stockquote-update
<SOAP-ENV:Envelope xmlns:SOAP-ENV="http://schemas.xmlsoap.org/soap/envelope/"
SOAP-ENV:encodingStyle="http://schemas.xmlsoap.org/soap/encoding/">
<SOAP-ENV:Body>
<m:buyStock xmlns:m="urn:stockquote:>
<symbol>IBM</symbol>
</m:buyStock>
</SOAP-ENV:Body>
</SOAP-ENV:Envelope>
```

The graphic illustrates how to invoke methods with the SOAP messages. In Web services terms, especially Web Service Definition Language (WSDL), a service is a collection of operations such as getQuote and buyStock. A body element namespace (urn:stockquote in our example) defines a service, and the name of the first body element indicates the operation.



The following is an example of WSDL for the getQuote operation:

```
<?xml version="1.0"?>
<definitions name="StockQuoteService-interface"
targetNamespace="http://www.getquote.com/StockQuoteService-interface"
xmlns:tns="http://www.getquote.com/StockQuoteService-interface"
xmlns:xsd="http://www.w3.org/2001/XMLSchema"
xmlns:soap="http://schemas.xmlsoap.org/wsdl/soap/"
xmlns="http://schemas.xmlsoap.org/wsdl/"
<message name="SymbolRequest">
<part name="return" type="xsd:string"/>
</message>
<portType name="StockQuoteService">
<operation name="getQuote">
<input message="tns:SymbolRequest"/>
<output message="tns:QuoteResponse"/>
</operation>
</portType>
<binding name="StockQuoteServiceBinding"
type="tns:StockQuoteService">
<soap:binding style="rpc" transport="http://schemas.xmlsoap.org/soap/http"/>
<operation name="getQuote">
<soap:operation soapAction="urn:stockquote-lookup"/>
<input>
<soap:body use="encoded" namespace="urn:stockquote"
encodingStyle="http://schemas.xmlsoap.org/soap/encoding"/>
</input>
<output>
<soap:body use="encoded" namespace="urn:stockquotes"
encodingStyle="http://schemas.xmlsoap.org/soap/encoding"/>
</output>
</operation>>
</binding>
</definition>
```

To build a set of cache policies for a Web services application configure WebSphere Application Server dynamic cache to recognize cacheable service operation of the operation.

WebSphere Application Server inspects the HTTP request to determine whether or not an incoming message can be cached based on the cache policies defined for an application. In this example, buyStock and stock-update are not cached, but stockquote-lookup is cached. In the cachespec.xml file for this Web application, the cache policies need defining for these services so that the dynamic cache can handle both SOAPAction and service operation.



WebSphere Application Server uses the operation and the message body in Web services cache IDs, each of which has a component associated with them. Therefore, each Web services <cache-id> rule contains only two components. The first is for the operation. Because you can perform the stockquote-lookup operation by either using a SOAPAction header or a service operation in the body, you must define two different <cache-id> elements, one for each method. The second component is of type "body", and defines how WebSphere Application Server should incorporate the message body into the cache ID. You can use a hash of the body, although it is legal to use the literal incoming message in the ID.

The incoming HTTP request is analyzed by WebSphere Application Server to determine which of the <cache-id> rules match. Then, the rules are applied to form cache or invalidation IDs.

The following is sample code of a cachespec.xml file defining SOAPAction and servicesOperation rules:

```
<cache>
<cache-entry>
<class>webservice</class>
<name>/soap/servlet/soaprouter</name>
<sharing-policy>not-shared</sharing-policy>
<cache-id>
<component id="" type=SOAPAction>
<value>urn:stockquote-lookup</value>
</component>
<component id="Hash" type="SOAPEnvelope"/>
<timeout>3600</timeout>
<priority>1</priority>
</cache-id>
<cache-id>
<component id="" type="serviceOperation">
<value>urn:stockquote:getQuote</value>
</component>
<component id="Hash" type="SOAPEnvelope"/>
<timeout>3600</timeout>
<priority>1</priority>
</cache-id>
</cache-entry>
</cache>
```

---

## Example: Configuring the dynamic cache

This example puts all the steps together for configuring the dynamic cache with the cachespec.xml file, showing the use of the cache ID generation rules, dependency IDs, and invalidation rules.

Suppose we have a servlet which is used to manage a simple news site. This servlet uses the query parameter "action" to determine whether the request is being used to "view" news or "update" news (used by the administrator). Further, another query parameter "category" is used to select the news category. Further, suppose that this site supports an optional customized layout, which is stored in the user's session using the attribute name "layout". Here are example URL requests to this servlet:

<http://yourhost/yourwebapp/newscontroller?action=view&category=sports> (Returns a news page for the sports category )

<http://yourhost/yourwebapp/newscontroller?action=view&category=money> (Returns a news page for the money category)

<http://yourhost/yourwebapp/newscontroller?action=update&category=fashion> (Allows the administrator to update news in the fashion category)

Here are the steps for configuring dynamic cache with cachespec.xml, using the information provided to you:

1. Define the cache-entry elements necessary to identify the servlet. In this case, the servlet's URI is "newscontroller" so this will be our cache-entry's name element. Also, since we are caching a servlet/JavaServer Page (JSP), the cache-entry class is "servlet".

```
<cache-entry>
<name> /newscontroller </name>
<class>servlet </class>
</cache-entry>
```

2. Define cache ID generation rules. For this servlet, we only want to cache when action=view, so one component of the cache ID will be the parameter "action" when the value equals "view". The news category is also an essential part of the cache ID. Finally, the optional session attribute for the user's layout is included in the cache ID. The cache-entry now looks like this:

```
<cache-entry>
<name> /newscontroller </name>
<class>servlet </class>
<cache-id>
<component id="action" type="parameter">
<value>view</value>
<required>true</required>
</component>
<component id="category" type="parameter">
<required>true</required>
</component>
<component id="layout" type="session">
<required>false</required>
</component>
</cache-id>
</cache-entry>
```

3. Define dependency ID rules. For this servlet, a dependency ID will be added for the category. Later, when the category is invalidated due to an update event, all views of that news category will be invalidated. After adding our dependency-id, the cache-entry now looks like this:

```
<cache-entry>
<name>newscontroller </name>
<class>servlet </class>
<cache-id>
<component id="action" type="parameter">
<value>view</value>
<required>true</required>
</component>
<component id="category" type="parameter">
<required>true</required>
</component>
<component id="layout" type="session">
<required>false</required>
</component>
</cache-id>
<dependency-id>category
<component id="category" type="parameter">
<required>true</required>
</component>
</dependency-id>
</cache-entry>
```

4. Define invalidation rules. Since we defined a category dependency ID, we will now define an invalidation rule to invalidate the category when action=update. To incorporate the conditional logic, we will add "ignore-value" components into the invalidation rule. These components will not add to the output of the invalidation ID, but will only determine whether or not the invalidation ID is created and executed. The final cache-entry now looks like this:

```
<cache-entry>
<name>newscontroller </name>
<class>servlet </class>
<cache-id>
<component id="action" type="parameter">
<value>view</value>
<required>true</required>
```

```

 </component>
 <component id="category" type="parameter">
 <required>true</required>
 </component>
 <component id="layout" type="session">
 <required>false</required>
 </component>
 </cache-id>
 <dependency-id>category
 <component id="category" type="parameter">
 <required>true</required>
 </component>
 </dependency-id>
 <invalidation>category
 <component id="action" type="parameter" ignore-value="true">
 <value>update</value>
 <required>true</required>
 </component>
 <component id="category" type="parameter">
 <required>true</required>
 </component>
 </invalidation>
</cache-entry>

```

---

## Cache monitor

Cache monitor is an installable Web application that provides a real-time view of the current state of dynamic cache. You use it to help verify that dynamic cache is operating as expected. The only way to manipulate the data in the cache is by using the cache monitor. It provides a GUI interface to manually change data.

Cache monitor does not provide a view for cache instance resources.

Cache monitor provides a way to:

- **Verify the configuration of dynamic cache**

The WebSphere Application Server administrative console provides ways to enable the dynamic cache service and configure properties, such as maximum size of the cache and disk offload location, as well as advanced features such as controlling external caches. Cache monitor offers a way for dynamic cache users to verify the configuration of the dynamic cache by providing a convenient view of the configured features and properties in the cache monitor.

- **Verify the cache policies**

To cache an object, WebSphere Application Server must know how to generate unique IDs for different invocations of that object. This is performed by providing rules for each cacheable object in the `cachespec.xml` file, found inside the Web module `WEB-INF` or enterprise bean `META-INF` directory. Each cacheable object can have multiple cache ID rules that execute in sequence until either a rule returns a cache ID or no more rules remain to execute. If none of the cache ID generation rules produce a valid cache ID, then the object is not cached. Since there can be multiple `cachespec.xml` files with multiple cache ID rules, cache monitor provides a convenient way to verify the policies of each object. It offers a view of all the cache policies currently loaded in dynamic cache. This view is also convenient to verify that the `cachespec.xml` file was read by the dynamic cache without errors.

- **Monitor cache statistics**

Cache monitor provides a view of the essential cache data, such as number of cache hits, cache misses, and number of entries in cache. This helps to tune the cache configuration optimally to get the best performance improvement out of dynamic cache. For example, if the number of used entries is often high, and entries are being removed and recreated, one might consider increasing the maximum size of the cache or enabling disk offload.

- **Monitor the data flowing through the cache**

Once a cacheable object is invoked, dynamic cache creates a cache entry for it that contains the output of the execution and metadata, such as time to live, sharing policy, etc. Entries are distinguished by a

unique ID string that is based on the rules specified in the `cachespec.xml` file for this objects name. Objects with the same name may generate multiple cache IDs for different invocations, based on request parameters and attributes for each invocation. Cache monitor provides a view of all the cache entries currently in cache, based on the unique ID. It also provides a view of the group of cache entries that share a common name (also known as template). Cache entries can also be grouped together by a dependency ID, which is used to invalidate the entire group of entries dependent on a common entity. Therefore, cache monitor also provides a view of the group of cache entries that share a common dependency ID.

For each entry, cache monitor also displays metadata, such as time to live, priority and sharing-policy, and provides a view of the output that has been cached. This helps the customer to verify which pages have been cached, that the pages have been cached with the right attributes such as time to live, priority, etc., and that the pages have the right content.

- **Monitor the data in the edge cache**

Dynamic cache provides support to recognize the presence of an Edge Side Include (ESI) processor and to generate ESI include tags and appropriate cache policies for edge cacheable fragments. The ESI processor has the ability to cache whole pages, as well as fragments, providing a higher cache hit ratio. There can be multiple ESI processors running on multiple hosts configured for caching.

Cache monitor provides a list of all ESI processes and their hosts that are enabled for caching. It also provides a way to select a host or a processor, and view its edge cache statistics as well as current cache entries.

- **View the data offloaded to the disk**

By default, when the number of cache entries reaches the configured limit for a given server, eviction of cache entries occurs, allowing new entries to enter the cache service. The dynamic cache includes the disk offload feature that copies the evicted cache entries to disk for future access. Cache monitor offers a view of the content offloaded to disk that corresponds to the view of contents cached in memory.

- **Manage the data in the cache**

Besides displaying cache content, cache monitor also provides some basic operations on the data in the cache:

- Removing an entry from the cache
- Removing all entries for a certain dependency ID
- Removing all entries for a certain name (template)
- Moving an entry to the front of the least recently used queue to avoid eviction
- Moving an entry from the disk to the cache
- Clearing the entire contents of the cache
- Clearing the contents of the disk cache

These functions are useful for dynamic cache customers, as they provide a way to manually change the state of the cache without having to restart the server.

## Edge cache statistics

Cache monitor provides a view of the edge cache statistics.

The following statistics are available:

- **ESI Processors.** Number of processes configured as edge caches.
- **Number of Edge Cached Entries.** Number of entries currently cached on all edge servers and processes.
- **Cache Hits.** Number of requests that match entries on edge servers.
- **Cache Misses By URL.** A cache policy does not exist on the edge server, for the requested template.

**Note:**

- The initial ESI request for a template that has a cache policy on WebSphere Application Server will result in a miss.
- Every request for a template that does not have a cache policy on WebSphere Application Server will result in a miss by URL on the edge server.

- **Cache Misses By Cache ID.** The policy for the requested template exists on the edge server, and a cache ID is created, based on the ID rules and the request attributes, but the cache entry for this ID does not exist.

**Note:** If the policy exists on the edge server for the requested template, but a cache ID match is not found, based on the ID rules and the request attributes, it is not treated as a cache miss.

- **Cache Timeouts.** Number of entries removed from the edge cache, based on the timeout value.
- **Evictions.** Number of entries removed from the edge cache, due to invalidations received from WebSphere Application Server.

---

## Troubleshooting the dynamic cache service

Complete the steps below to resolve problems that you think are related to the dynamic cache service.

1. Review the JVM logs for your application server. Messages prefaced with *DYNA* result from dynamic cache service operations.
  - a. View the JVM logs for your application server. Each server has its own JVM log file. For example, if your server is named *Member\_1*, the JVM log is located in the subdirectory *install\_root/logs/Member\_1*. To use the administration console to review the JVM logs, click **Troubleshooting > Logs and Trace > server\_name > JVM Logs > Runtime > View**.
  - b. Find any messages prefaced with *DYNA* in the JVM logs, and write down the message IDs. A sample message having the message ID *DYNA0030E* follows:  
 DYNA0030E: "property" element is missing required attribute "name".
  - c. Find the message for each message ID in the information center for WebSphere Application Server. In the information center navigation tree, click **product\_name > Reference > Messages > DYNA** to view dynamic cache service messages.
  - d. Read the message **Explanation** and **User Action** statements. A search for the message ID *DYNA0030E* displays a page having the following message:  
**DYNA0030E: "property" element is missing required attribute "name".**  
**Explanation:** A required attribute was missing in the cache configuration.  
**User Action:** Add the required attribute to your cache configuration file.  
 This explanation and user action suggests that you can fix the problem by adding or correcting a required attribute in the cache configuration file.
  - e. Try the solutions stated under **User Action** in the *DYNA* messages.
2. Use the cache monitor to determine whether the dynamic cache service is functioning as expected. The cache monitor is an installable Web application that displays simple cache statistics, cache entries, and cache policy information.
3. If you have completed the preceding steps and still cannot resolve the problem, contact your IBM software support representative. Use the collector tool (*collector.bat* or *collector.sh* located in the *bin* directory) to gather trace information and other configuration information for the support team to diagnose the problem. The collector tool gathers dynamic cache service files and packages them into a JAR file. The IBM representative can specify when and where to send the JAR file. The IBM representative might ask you to complete a diagnostic trace. To enable tracing in the administrative console, click **Troubleshooting > Logs and Trace > server\_name > Diagnostic Trace** and specify **Enable trace with the following specification**. The IBM representative can tell you what trace specification to enter. Note that dynamic cache trace files can become large in a short period of time; you can limit the size of the trace file by starting the trace, immediately recreating the problem, and immediately stopping the trace.

For current information available from IBM Support on known problems and their resolution, see the IBM Support page.

For technical support on dynamic cache service, see the IBM Support page.

## Troubleshooting tips for the dynamic cache service

The dynamic cache service works within an application server Java virtual machine (JVM), intercepting calls to cacheable objects. This article describes some common run-time and configuration problems and remedies.

### Servlets are not cached

**Recommended response** Enable servlet caching. On the Web container page of the administrative console, select the **Enable servlet caching** check box.

### Cache entries are not written to disk

**Explanation** Cache entries are written to disk when the cache is full and new entries are added to the memory cache. Cache entries also are written to disk when the `flushToDiskOnStop` system property is set and the server is stopped.

**Recommended response** Verify that **Disk offload** is enabled on the Dynamic Cache Service page of the administrative console. Also verify that cache entries written to disk are serializable and do not have the `PersistToDisk` configuration set to `false`.

### Some servlets are not replicated or written to disk

**Recommended response** Ensure that the attributes and response are serializable. If you do not want to store the attributes, use the following property in your cache policy:  
`<property name="save-attributes">false</property>`

### Dynamic cache service does not cache fragments on the Edge

**Recommended response** Set the `EdgeCacheable` property to `true` in the cache policy for those entries that are to be cached on the Edge.  
`<property name="EdgeCacheable">true</property>`

### Dynamic cache invalidations are not sent to the IBM HTTP Server (IHS) plug-in

**Explanation** The `DynaCacheEsi`.ear file is required to send invalidations to external caches.  
**Recommended response** Install `DynaCacheEsi`.ear using the administrative console.

### Cache entries are evicted often

**Problem** The cache is full and new entries are added to the cache.  
**Explanation** Cache entries are evicted when the cache is full and new entries are added to the cache. A LRU eviction mechanism removes the least recently used entry to make space for the new entries.  
**Recommended response** Either enable **Disk offload** on the Dynamic Cache Service page of the administrative console so the entries are written to disk. Or, increase the cache size to accommodate more entries in the cache.

### Cache entries in disk with timeout set to 0 expire after one day

<b>Explanation</b>	The maximum lifetime of an entry in disk cache is 24 hours. A timeout of 0 in the cache policy configures these entries to stay in disk cache for one whole day, unless they are evicted earlier.
<b>Recommended response</b>	Set the timeout for the cache policy to a number greater than 0.

### I cannot monitor cache entries on the Edge

<b>Explanation</b>	Use the cache monitor for monitoring contents in memory cache, disk cache and external caches (Edge cache). For the ESI processor's cache to be visible in the cache monitor, the DynaCacheEsi.ear application must be installed and the esiInvalidationMonitor property must be set to true in the plugin-cfg.xml file.
<b>Recommended response</b>	Install the DynaCacheEsi.ear application and set the esiInvalidationMonitor property to true in the plugin-cfg.xml file.

### I want to cache static contents using the dynamic cache service

<b>Explanation</b>	You can cache static contents using the dynamic cache service. Static contents in WebSphere application server are served by the SimpleFileServlet file.
<b>Recommended response</b>	Create a cache policy for the class <code>com.ibm.ws.webcontainer.servlet.SimpleFileServlet.class</code> to cache static contents. It is advisable to use the dynamic cache service for caching more expensive dynamic contents than static contents.

### I want to tune cache for my environment

<b>Recommended response</b>	Use the Tivoli Performance viewer to study the caching behavior for your applications. Also, do the following: <ul style="list-style-type: none"><li>• Increase the priority of cache entries that are expensive to regenerate.</li><li>• Modify timeout of entries so that they stay in memory as long as they are valid.</li><li>• Enable disk offload to store LRU evicted entries.</li><li>• Increase the cache size.</li></ul>
-----------------------------	-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------





---

## Chapter 28. Using user profiles

**Note:** User Profile Manager API is deprecated in the current release, and there is no replacement available.

IBM WebSphere Application Server provides a service for processing user profiles, called the *User Profile Manager*. The service is provided in the form of an EJB entity bean that servlets can call whenever they are required to access a user profile.

The key activities for implementing user profiles are summarized.

1. Customize the user profile support as necessary. Options include:
  - Using the data representation class with the name-value pairs it currently supports (no action required)
  - Extending the data representation class to support additional, arbitrary name-value pairs
  - Adding columns to the base user profile representation
  - Extending the User Profile enterprise bean to import existing databases

Evaluate whether the user profile representation provided by IBM represents the kind of data you want to keep about your users. You might find it desirable to customize the IBM user profile support.

2. Create or modify servlets to use the User Profile Manager and related user profile support classes to maintain user profiles on behalf of Web applications.
3. Assemble your application.
4. Deploy your application.
5. Ensure the administrator appropriately configures User Profile Managers using `userprofile.xml` file. If the programmer and administrator are not the same person, the programmer might need to provide settings information to the administrator, based on how the programmer implemented user profiles.

---

### User profile

Some applications collect data about the users with which they interact. The data is stored in a database. The next time the user interacts with the application, the application recalls the data.

Because the application already knows something about the user, it can provide the user with a more personalized experience.

User profiles provide a means by which a company can maintain and manage database tables containing fields for demographic data and use those tables to interact with a database of individual customers or other users on the company system.

For example, when a repeat user logs onto a Web site that supports user profiles, the Web site can display headlines and advertising tailored to the shopping preferences of that user. The site can address the user by logon name.

An application implementing user profiles requires database access for storing the user profile data it gathers.

---

### UserProfileManager class

Servlets and other application building blocks requiring user profile support should make calls to the class:  
`com.ibm.websphere.userprofile.UserProfileManager`

The class supports the following functions:

- Creating and deleting user profiles
- Getting and updating (cached and immediate) to and from the database

- Getting user profiles for read-only tasks
- Performing queries on database columns

---

## User profile development options

The application developer has a few options for customizing the user profile support provided by IBM WebSphere Application Server. The Related information provides instructions and additional details about each option.

### Extending the data represented in user profiles

Web applications can maintain several pieces of data about users. You can extend the data representation to allow the collection of arbitrary name-value pairs.

Use the following interface with the `com.ibm.websphere.userprofile.UserProfileExtender` class to extend a user profile hash table:

```
com.ibm.websphere.userprofile.UserProfileProperties
```

This action enables you to place arbitrary name-value pairs in the user profile. Extending the hash table is similar to using the `java.util.Dictionary` class in the base JDK 1.x, or any of the classes that extend it.

### Adding columns to the base user profile implementation

Application developers can customize user profiles by adding columns to the base user profile implementation. Adding new columns is accomplished by implementing the interface:

```
com.ibm.websphere.userprofile.UserProfileExtender
```

and extending the base class:

```
com.ibm.servlet.personalization.userprofile.UserProfile
```

The application developer can add columns to but not delete columns from the base implementation.

Adding columns is a two-step process, as follows:

1. Extend the `UserProfile` class.
2. Modify your existing servlets to use the new columns.

Several examples are available to demonstrate how to extend the base user profile implementation and utilize the extension with a servlet.

Example	Description
<code>UPServletExample.java</code>	Demonstrates how a servlet opens a user profile and prints the fields contained within.
<code>UserProfileExtendedSample.java</code>	Shows how to extend the <code>UserProfile</code> class to add a column to the user profile for a cellular phone number.  The WebSphere Application Server administrator configures the User Profile Manager to point to the extended class.
<code>UPServletExampleExtended.java</code>	Shows how to modify the <code>UPServletExample</code> servlet to include the cellular phone number in the output.
<code>UserProfileExtended.java</code>	Shows how to extend a hash table to place arbitrary name-value pairs into the user profile.
<code>UPServletExtended.java</code>	Shows how to extend the servlet. When any of the newly added columns are removed or replaced, look for the table named "USERPROFILE" in the database to which the user profile is configured and drop that table.

The examples are encoded in HTML for viewing in a browser. The documentation directory also contains nonHTML versions (.java files) that are ready for use.

## Extending the User Profile enterprise bean and importing legacy databases

Application developers can extend the User Profile enterprise bean itself and import legacy databases into the user profile. The main advantage in extending the User Profile enterprise bean is to gain the ability to import existing databases into the user profile. You can also extend this enterprise bean to add columns to the base user profile implementation.

### Example: UPServletExample.java

```
import java.io.*;
import java.util.*;

import javax.servlet.*;
import javax.servlet.http.*;

import com.ibm.websphere.userprofile.UserProfile;
import com.ibm.websphere.userprofile.UserProfileManager;
import com.ibm.websphere.userprofile.UserProfileCreateException;
import com.ibm.websphere.userprofile.UserProfileFinderException;
import com.ibm.websphere.userprofile.UserProfileRemoveException;

//Creates a Userprofile using the new API

public class UPServlet_ReadWrite extends HttpServlet {

 public void doGet (HttpServletRequest req, HttpServletResponse res)
 throws ServletException, IOException {

 PrintWriter out;
 res.setContentType("text/html");
 out = res.getWriter();

 UserProfileManager manager = UserProfileManager.getUserProfileManager();
 UserProfile userprofile;

 try {

 //Try creating the UserProfile
 userprofile = manager.addUserProfile("bpink");

 } catch(UserProfileCreateException e1) {

 try {

 //Try finding the existing in readWrite mode.
 //Second argument indicates whether we want to get userprofile
 //in read only mode or read write mode.

 userprofile = manager.getUserProfile("bpink",true);

 } catch(UserProfileFinderException e) {
 e.printStackTrace();
 return;
 }

 }

 }

}
```

```

//Set the properties

userprofile.setAddress1("myaddress1");
userprofile.setAddress2("myaddress2");
userprofile.setFirstName("Pinkowski");
userprofile.setSurName("Ben");
userprofile.setDayPhone("555-6677");
userprofile.setNightPhone("556-6765");
userprofile.setCity("MYCITY");
userprofile.setNation("myCountry");
userprofile.setEmployer("MyEmployer");
userprofile.setFax("7823470");

userprofile.setLanguage("mylanguage");
userprofile.setEmail("MyEmail@email");
userprofile.setStateOrProvince("myState");
userprofile.setPostalCode("xxxxx");

//Freeing resources held by userprofile
manager.releaseResources(userprofile);
userprofile=null;

//Checking whether it updated the info

try {

 //Getting the existing userprofile in ReadOnly mode.

 userprofile = manager.getUserProfile("bpink",false);

} catch(UserProfileFinderException e1) {

 out.println("Error finding ");
 e1.printStackTrace();
 return;
}

//Displaying the properties of userprofile

out.println(userprofile.getAddress1()+"
");
out.println(userprofile.getAddress2()+"
");;
out.println(userprofile.getFirstName()+"
");;
out.println(userprofile.getSurName()+"
");;
out.println(userprofile.getDayPhone()+"
");;
out.println(userprofile.getNightPhone()+"
");;
out.println(userprofile.getCity()+"
");;
out.println(userprofile.getNation()+"
");;
out.println(userprofile.getEmployer()+"
");;
out.println(userprofile.getFax()+"
");;
out.println(userprofile.getLanguage()+"
");;
out.println(userprofile.getEmail()+"
");;
out.println(userprofile.getStateOrProvince()+"
");;
out.println(userprofile.getPostalCode()+"
");;

//Freeing resources held by userprofile
manager.releaseResources(userprofile);
}
}

```

## Example: UserProfileExtendedSample.java

```

/* -----
** Copyright 1997-99 IBM Corporation. All rights reserved.
**

```

```

** -----
*/
package com.ibm.servlet.personalization.userprofile;

import com.ibm.servlet.personalization.userprofile.UserProfile;
import com.ibm.websphere.userprofile.UserProfileExtender;

//Extensions of UserProfile to add new Columns should implement UserProfileExtender
public class UserProfileExtendedSample
 extends com.ibm.servlet.personalization.userprofile.UserProfile
 implements UserProfileExtender {

 //New column that is being added by this
 //derived class.
 public String cellPhone;

 //Manager Class will call this method to append new Column types.
 //If UserProfile class is extended to append new columns
 //TOTAL COLUMNS: Base Class columns + columns returned by this class

 public String[] getNewColumns() {
 //If variable name is "cellPhone," you need to
 //return "cellPhone" in array format. JDBC equivalent will be
 //generated automatically. You can add multiple columns.
 //For multiple columns: String newCol={"fieldName1","fieldName2",...};
 String[] newCol={"cellPhone"};
 return newCol;
 }

 public String getCellPhone() {
 // Need to call this method to
 // get the things from persistence store.
 return(String)getByType("cellPhone");
 }

 public void setCellPhone(String value) {
 cellPhone = value;
 //Call this method to store the
 //things in persistence store
 setByType("cellPhone", value);
 }
}

```

## Example: UPServletExampleExtended.java

```

import java.io.*;
import java.util.*;

import javax.servlet.*;
import javax.servlet.http.*;

import com.ibm.websphere.userprofile.UserProfile;
import com.ibm.websphere.userprofile.UserProfileManager;
import com.ibm.websphere.userprofile.UserProfileCreateException;
import com.ibm.websphere.userprofile.UserProfileFinderException;
import com.ibm.websphere.userprofile.UserProfileRemoveException;
import com.ibm.servlet.personalization.userprofile.UserProfileExtendedSample;

public class UPServletExtendedSample extends HttpServlet {

 public void doGet (HttpServletRequest req, HttpServletResponse res)
 throws ServletException, IOException
 {

```

```

UserProfileManager manager = UserProfileManager.getUserProfileManager();
UserProfile userprofile;

PrintWriter out;

res.setContentType("text/html");
out = res.getWriter();

try {
 //try Creating the UserProfile
 userprofile = manager.addUserProfile("bpink");
} catch(UserProfileCreateException e1) {
 try { //try finding the existing in readWrite mode
 userprofile = manager.getUserProfile("bpink",true);
 } catch(UserProfileFinderException e) {
 e.printStackTrace();
 return;
 }
}

userprofile.setAddress1("myaddress1");
userprofile.setAddress2("myaddress2");
userprofile.setFirstName("Pinkowski");
userprofile.setSurName("Ben");
userprofile.setDayPhone("555-6677");
userprofile.setNightPhone("556-6765");
userprofile.setCity("MYCITY");
userprofile.setNation("myCountry");
userprofile.setEmployer("MyEmployer");
userprofile.setFax("7823470");
userprofile.setLanguage("mylanguage");
userprofile.setEmail("MyEmail@email");
userprofile.setStateOrProvince("myState");
userprofile.setPostalCode("xxxxx");

//calling setCellPhone
((com.ibm.servlet.personalization.userprofile.UserProfileExtendedSample)
 userprofile).setCellPhone("346-4588");

//Freeing resources held by userprofile
manager.releaseResources(userprofile);
userprofile=null;

//Checking whether it updated the info
try {
 //Getting the existing userprofile
 userprofile = manager.getUserProfile("bpink",false);
} catch(UserProfileFinderException e1) {
 out.println("Error finding ");
 e1.printStackTrace();
 return;
}

```



```

 }

 out.println(userprofile.getAddress1()+"
");
 out.println(userprofile.getAddress2()+"
");;
 out.println(userprofile.getFirstName()+"
");;
 out.println(userprofile.getSurName()+"
");
 out.println(userprofile.getDayPhone()+"
");;
 out.println(userprofile.getNightPhone()+"
");;
 out.println(userprofile.getCity()+"
");
 out.println(userprofile.getNation()+"
");;
 out.println(userprofile.getEmployer()+"
");;
 out.println(userprofile.getFax()+"
");;
 out.println(userprofile.getLanguage()+"
");;
 out.println(userprofile.getEmail()+"
");;
 out.println(userprofile.getStateOrProvince()+"
");;
 out.println(userprofile.getPostalCode()+"
");

 //Calling getCellPhone
 out.println(((UserProfileExtendedSample)userprofile).getCellPhone()+"
");

 //Freeing resources held by userprofile
 manager.releaseResources(userprofile);
 userprofile=null;

 //For getting values by cellPhone
 out.println("

Retreiving by Cell Phone
");
 Enumeration enum = manager.findUserProfiles("cellPhone","346-4588");
 while(enum.hasMoreElements()) {

 com.ibm.websphere.userprofile.UserProfile up =
 (com.ibm.websphere.userprofile.UserProfile)enum.nextElement();
 out.println("first name :"+up.getFirstName()+"
");

 //Freeing resources held by userprofile
 manager.releaseResources(up);
 }
}
}
}

```

## Example: UserProfileExtended.java

```

package com.ibm.servlet.personalization.userprofile;
/* -----
** Copyright 1997-99 IBM Corporation. All rights reserved.
** -----
*/
import java.util.*;

import com.ibm.servlet.personalization.userprofile.UserProfile;

import com.ibm.websphere.userprofile.UserProfileExtender;
import com.ibm.websphere.userprofile.UserProfileProperties;

public class UserProfileExtended extends UserProfile implements UserProfileExtender,
 UserProfileProperties {
 //New column that is being added by this
 //derived class.
 public Hashtable properties;

 static String propCol ="properties";

 //Manager Class will call this method to append new Column types
 //to SQL Strings. If UserProfile class is extended to append new columns

```

```

//it should implement UserProfileExtender.
//COLUMNS: Base Class columns + columns returned by this class

public String[] getNewColumns() {
 //if variable name is properties, you need to
 //return "properties" . JDBC equivalent will be
 //generated automatically.
 String[] newCol={propCol};
 return newCol;
}

public Object getValue(String key) {
 // Need to call this method to
 // get the things from persistent store
 properties = (Hashtable) getByType(propCol);

 if(properties != null)
 return properties.get(key);

 else return null;
}

public void putValue(String key, Object value) {
 properties =(Hashtable) getByType(propCol);

 if(properties == null)
 properties = new Hashtable();

 properties.put(key,value);

 //store in persistent store
 setByType(propCol, properties);
}

public void removeValue(String key) {
 properties = (Hashtable) getByType(propCol);

 if(properties == null)
 return;

 properties.remove(key);

 //store in persistent store
 setByType(propCol, properties);
}
}

```

## Example: UPServletExtended.java

```

import java.io.*;
import java.util.*;

import javax.servlet.*;
import javax.servlet.http.*;

import com.ibm.websphere.userprofile.UserProfile;
import com.ibm.websphere.userprofile.UserProfileManager;
import com.ibm.websphere.userprofile.UserProfileCreateException;
import com.ibm.websphere.userprofile.UserProfileFinderException;
import com.ibm.websphere.userprofile.UserProfileRemoveException;
import com.ibm.websphere.userprofile.UserProfileProperties;

public class UPServletExtended extends HttpServlet {

```

```

public void doGet (HttpServletRequest req, HttpServletResponse res)
throws ServletException, IOException
{
 UserProfileManager manager = UserProfileManager.getUserProfileManager();
 UserProfile userprofile;

 PrintWriter out;

 res.setContentType("text/html");
 out = res.getWriter();

 try {

 //try Creating the UserProfile

 userprofile = manager.addUserProfile("bpink");

 } catch(UserProfileCreateException e1) {

 try { //try finding the existing in readWrite mode

 userprofile = manager.getUserProfile("bpink",true);
 } catch(UserProfileFinderException e) {
 e.printStackTrace();
 return;
 }

 }

 userprofile.setAddress1("myaddress1");
 userprofile.setAddress2("myaddress2");
 userprofile.setFirstName("Pinkowski");
 userprofile.setSurName("Ben");
 userprofile.setDayPhone("555-6677");
 userprofile.setNightPhone("556-6765");
 userprofile.setCity("MYCITY");
 userprofile.setNation("myCountry");
 userprofile.setEmployer("MyEmployer");
 userprofile.setFax("7823470");
 userprofile.setLanguage("mylanguage");
 userprofile.setEmail("MyEmail@email");
 userprofile.setStateOrProvince("myState");
 userprofile.setPostalCode("xxxxx");

 //calling putValue

 ((UserProfileProperties)userprofile).putValue("name","HHHHHHH");
 ((UserProfileProperties)userprofile).putValue("Date",new java.util.Date());

 //Freeing resources held by userprofile
 manager.releaseResources(userprofile);
 userprofile=null;

 //Checking whether it updated the info

 try {

 //Getting the existing userprofile

 userprofile = manager.getUserProfile("bpink",false);

 out.println(userprofile.getAddress1()+"
");
 out.println(userprofile.getAddress2()+"
");
 out.println(userprofile.getFirstName()+"
");
 out.println(userprofile.getSurName()+"
");
 out.println(userprofile.getDayPhone()+"
");
 }
}

```

```

 out.println(userprofile.getNightPhone()+"
");
 out.println(userprofile.getCity()+"
");
 out.println(userprofile.getNation()+"
");
 out.println(userprofile.getEmployer()+"
");
 out.println(userprofile.getFax()+"
");
 out.println(userprofile.getLanguage()+"
");
 out.println(userprofile.getEmail()+"
");
 out.println(userprofile.getStateOrProvince()+"
");
 out.println(userprofile.getPostalCode()+"
");

 //Getting the values

 out.println(((UserProfileProperties)userprofile).getValue("name")+"
");
 out.println(((UserProfileProperties)userprofile).getValue("Date")+"
");
 out.println("Removing Values ");
 ((UserProfileProperties)userprofile).removeValue("name");
 ((UserProfileProperties)userprofile).removeValue("Date");
 out.println(((UserProfileProperties)userprofile).getValue("name")+"
");
 out.println(((UserProfileProperties)userprofile).getValue("Date")+"
");

 //Freeing resources held by userprofile
 manager.releaseResources(userprofile);

 } catch(UserProfileFinderException e1) {

 out.println("Error finding ");
 e1.printStackTrace();
 return;
 }
}
}
}

```

---

## userprofile.xml

To *installation\_root/properties*, add a file named `userprofile.xml` in the following format. Specify enterprise bean class names; data wrapper class name; and JNDI names for the read-only bean, read/write bean, and data source (from step 2). You must also add user ID and password information for the JNDI data source

The following example file contains class names as provided in WebSphere Application Server. If data wrapper and enterprise bean class names are extended programmatically, change them accordingly.

```

<?xml version="1.0"?>

<userprofile>
 <userprofile-enabled>true</userprofile-enabled>
 <userprofile-wrapper-class>
 <classname>
 com.ibm.servlet.personalization.userprofile.UserProfile
 </classname>
 </userprofile-wrapper-class>
 <userprofile-manager-name>
 User Profile Manager
 </userprofile-manager-name>
 <userprofile-bean>
 <readonly-interface>
 com.ibm.servlet.personalization.userprofile.UP_ReadOnly
 </readonly-interface>
 <readwrite-interface>
 com.ibm.servlet.personalization.userprofile.UP_ReadWrite
 </readwrite-interface>
 <readonlyhome-interface>
 com.ibm.servlet.personalization.userprofile.UP_ReadOnlyHome
 </readonlyhome-interface>
 <readwritehome-interface>

```

```
 com.ibm.servlet.personalization.userprofile.UP_ReadWriteHome
 </readwritehome-interface>
 <readonly-JNDI-lookupName>UP_ReadOnlyHome</readonly-JNDI-lookupName>
<readwrite-JNDI-lookupName>UP_ReadWriteHome</readwrite-JNDI-lookupName>
 </userprofile-bean>

 <userprofile-store>
 <database-userid></database-userid>
 <database-password></database-password>
 <database-datasource></database-datasource>
 </userprofile-store>

</userprofile>
```



---

## Chapter 29. Assembling applications with the Assembly Toolkit

Assemble enterprise application modules (EAR files) from new or existing Java 2 Platform, Enterprise Edition (J2EE) Version 1.2 or 1.3 modules, including these archives: Web application archives (WAR), resource adapter archives (RAR), enterprise bean (EJB) JAR files, and application client archives (JAR). This packaging and configuration of code artifacts into application modules or stand-alone Web modules is necessary for deploying the applications onto the application server.

The Assembly Toolkit replaces the Application Assembly Tool (AAT). The Assembly Toolkit consists of the J2EE Perspective of the WebSphere Studio Application Developer product. With the Assembly Toolkit, you can create and modify J2EE applications and modules, edit deployment descriptors, and map databases.

The Assembly Toolkit is one of the tools provided by the Application Server Toolkit (ASTK). Follow instructions available with the ASTK to install the Assembly Toolkit.

Gather the code artifacts that you want to package into one or more assembled modules. Code artifacts include these items that you have created and unit tested in your favorite integrated development environment:

- Enterprise beans
- Servlets, JavaServer Pages (JSP) files and other Web components
- Resource adapter (*connector*) implementations
- Application clients
- Other supporting classes and files

The Assembly Toolkit provides extensive online documentation. The articles on Assembly Toolkit provided in the Information Center supplement that documentation.

1. Start the Assembly Toolkit.
2. Optional: Read the online documentation for the Assembly Toolkit.
  - Read the section **Assembly Tool** on the Welcome to the Application Server Toolkit page. To access this page, click **Help > Welcome > Application Server Toolkit**.
  - Click **Help > Help Contents > Assembly Toolkit information**. The displayed documentation provides extensive information about the Assembly Toolkit.
  - Press F1 to access information specific to an Assembly Toolkit view or window.
  - Visit the information center for IBM WebSphere Studio Application Developer at <http://publib.boulder.ibm.com/infocenter/wsphelp/index.jsp>. Click **WebSphere Studio Application Developer > J2EE development**. The documentation in the Information Center for WebSphere Studio is similar to that in the Assembly Toolkit online information.
  - See the article “Assembly Toolkit: Resources for learning” on page 1407” for additional sources.
3. Optional: Open the J2EE perspective to work with J2EE projects. Click **Window > Open Perspective > Other > J2EE**.
4. Optional: Open the J2EE Hierarchy view. Click **Window > Show View > J2EE Hierarchy**. Other helpful views include the Project Navigator view (**Window > Show View > Other > J2EE > Project Navigator**) and the Navigator view (**Window > Show View > Navigator**).
5. **5.1+** Optional: To use an application assembly service of WebSphere Business Integration (WBI) Server Foundation, select the **Integration Server** target server. Available WBI assembly services include ActivitySession, Application profiling, Extended access, Extended messaging, Internationalization, Last participant support, and Pushdown.
  - a. Click **Project > Properties > J2EE**.
  - b. For **Target server**, select **Integration Server**. For this release, **Integration Server v5.1**.
  - c. Click **OK**.



6. Migrate EAR, WAR, enterprise bean JAR files, application client JAR files, or resource adapter RAR files created with the Application Assembly Tool (AAT) or a different tool to the Assembly Toolkit. To migrate files, import the files to the Assembly Toolkit.
7. Optional: Migrate a project from J2EE 1.2 to J2EE 1.3 using the J2EE Migration wizard. As part of the migration, you can migrate CMP 1.x beans to CMP 2.x beans. The J2EE Migration wizard is similar to the earconvert batch utility or the **File > ConvertEar** option of the AAT.
  - a. In the J2EE Hierarchy view, right-click the enterprise application project (EAR file) you want to migrate.
  - b. Click **Migrate > J2EE Migration Wizard**.
  - c. Follow the instructions in the wizard.
8. Create an enterprise application project to which you can add archive files. You can create an enterprise application project separately or when you create archive files such as the following:
  - Create a Web project.
  - Create an application client.
  - Create an enterprise bean (EJB) project.
  - Create a resource adapter (connector) project.
9. Edit the deployment descriptors as needed. You can edit deployment descriptors for enterprise application, Web, application client, and enterprise bean (EJB) modules.
10. Optional: Generate enterprise bean (EJB) to relational database (RDB) mappings for EJB modules.
11. Verify the archive files.
12. Generate code for deployment for EJB modules or for enterprise applications that use EJB modules.
13. Generate code for deployment for Web services-enabled modules or for enterprise applications that use Web service modules.
14. Optional: Test your completed module on a WebSphere Application Server installation. Right-click a module, click **Run on Server**, and follow the instructions in the displayed wizard. Note that **Run on Server** works on the Windows, Linux/Intel, and AIX operating systems only; you cannot deploy remotely from the Assembly Toolkit to a WebSphere Application Server installation on a UNIX operating system such as Solaris.

### Important

**Important:** Use **Run On Server** for unit testing only. Application Server Toolkit controls the WebSphere Application Server installation and, when an application is published remotely, the Toolkit overwrites the server configuration file for that server. Do not use on production servers.

For instructions on remote testing, see the article "Setting Up a Remote WebSphere Application Server in WebSphere Studio V5" at [http://www7b.boulder.ibm.com/wsdd/techjournal/0303\\_yuen/yuen.html](http://www7b.boulder.ibm.com/wsdd/techjournal/0303_yuen/yuen.html).

After assembling your applications, use a systems management tool to deploy the EAR or WAR files onto the application server. The systems management tool follows the security and deployment instructions defined in the deployment descriptor, and enables you to modify bindings specified within the Assembly Toolkit. The tool locates the required external resources that the application uses, such as enterprise beans and databases.

Select a tool to use:

- Administrative console installation pages (GUI)
- Java administrative programs (programming). See "Deploying and managing using programming" in the Information Center.
- wsadmin AdminApp install command (scripting). See "Deploying and managing using scripting" in the Information Center.

If you are uncertain of which systems management tool to use, try using the administrative console.

If your application has a large number of modules, it might not install successfully onto a server. Package your application so that the .ear file contains necessary modules only. Modules can include metadata for the modules such as information on deployment descriptors, bindings, and IBM extensions.

Use the administrative console at installation to complete the security instructions defined in the deployment descriptor and to locate required external resources, such as enterprise beans and databases. You can add configuration properties and redefine binding properties defined in the Assembly Toolkit.

---

## Application assembly and J2EE applications

Application assembly is the process of creating an enterprise archive (EAR) file containing all files related to an application, as well as an XML deployment descriptor for the application. This configuration and packaging prepares the application for deployment onto an application server.

EAR files are comprised of the following archives:

- Enterprise bean JAR files (known as EJB modules)
- Web archive (WAR) files (known as Web modules)
- Application client JAR files (known as client modules)
- Resource adapter archive (RAR) files (known as resource adapter modules)

Ensure that modules are contained in an EAR file so that they can be deployed onto the server. The exceptions are WAR modules, which you can deploy individually. Although WAR modules can contain regular JAR files, they cannot contain the other module types described previously.

The assembly process includes the following actions:

- Selecting all of the files to include in the module.
- Creating a deployment descriptor containing instructions for module deployment on the application server.

As you configure properties using the Assembly Toolkit, the tool generates the deployment descriptor for you. While the Assembly Toolkit graphical interface is recommended, you can also edit descriptors directly in your favorite XML editor.

- Packaging modules into a single EAR file, which contains one or more files in a compressed format.

---

## Archive support in Version 5

The following archives and Web components are supported in Version 5.0:

- Java 2 Platform, Enterprise Edition (J2EE) Version 1.2 and 1.3 enterprise archive (EAR) files
- Enterprise bean (EJB) 2.0 JAR files
- Servlet 2.3 Web archive (WAR) files
- Application client 1.2 and 1.3 JAR files
- Connector 1.0 resource adapter archive (RAR) files

These archive files and Web components are back-level and can be read but not created or changed:

- J2EE 1.2 EAR files
- EJB 1.1 JAR files
- Servlet 2.2 WAR files
- Application client 1.2 JAR files

---

## Starting the Assembly Toolkit

The Assembly Toolkit provides a graphical interface for packaging code artifacts into various archives (modules) and configuring related Java 2 Platform, Enterprise Edition (J2EE) Version 1.2 or 1.3 compliant deployment descriptors.

The Assembly Toolkit is a component of the Application Server Toolkit (ASTK). To install the Assembly Toolkit, follow the installation instructions for the ASTK and, when prompted by the ASTK installation program, select to install the application assembly toolkit.

If you have installed the Assembly Toolkit component of ASTK previously and you install the Assembly Toolkit again, you must delete the workspace of the previous installation of ASTK before starting the Assembly Toolkit. The default workspace directory is *my\_directory\IBM\astk\workspace*. If you do not delete the workspace for a previous installation of ASTK, you might encounter error messages such as the following when starting the Assembly Toolkit:

Problems during startup. Check the ".log" file in the ".metadata" directory of your workspace.

1. Run the `astk` executable.
2. In the Application Server Toolkit window, specify the workspace directory and click **OK** to launch the graphical interface.

The navigation tree displays a hierarchical structure used to build the contents of a new module, or to work with the contents of an existing module.

Consider whether you have any existing J2EE 1.2 application modules that you would like to migrate to J2EE 1.3.

You can import or create new modules of the following types, to assemble into an application module later:

- Assembling enterprise bean (EJB) modules
- Assembling Web modules
- Assembling application client modules
- Assembling resource adapter modules

Rather than import or create new modules to assemble an application, you can proceed directly to assembling a new application module. While assembling an application module, you can create any new modules that you need.

## astk command

The **astk** command starts the Application Server Toolkit. The command file is the `astk` executable file.

### Location of the command file

The `astk` executable file resides in the main installation directory for Application Server Toolkit (ASTK).

### Command syntax

- Issue the **astk** command:

```
ASTK_install_root/astk
```

Or, double-click the **ASTK** icon.

- The **astk** command has no options or command-line parameters.

### Usage notes

- Is the `astk` executable file read-only?  
Yes
- Is this file updated by a product component?  
No
- How and when are the contents of this file used?

The `astk` executable file provides the Assembly Toolkit component for the Application Server Toolkit. Run the `astk` executable file to start the Assembly Toolkit. There are no parameters or command-line options.

---

## Migrating code artifacts to the Assembly Toolkit

You can migrate enterprise archive (EAR), Web archive (WAR), enterprise bean JAR, application client JAR, resource adapter archive (RAR) files created with the Application Assembly Tool (AAT) or a different tool to the Assembly Toolkit. To migrate files, import the files to the Assembly Toolkit.

1. Start the Assembly Toolkit.
2. Import an enterprise application.
3. Import a WAR file.
4. Import an application client file.
5. Import an enterprise bean JAR file.
6. Import a resource adapter RAR file. RAR files are also known as *connectors*.
7. Verify the archive files.
8. Generate code for deployment.

## Importing enterprise applications

You can import an enterprise archive (EAR) file and define a new enterprise application project using the Assembly Toolkit.

1. Start the Assembly Toolkit.
2. Click **File > Import**. Alternatively, you can right-click **Enterprise Applications** in a view such as the J2EE Hierarchy view and click **Import**. Or, on Windows platforms, you can drag the EAR file and drop it on a view.
3. In the Import dialog, specify the EAR file to import and the project name:
  - a. Click **EAR file > Next**.
  - b. Specify the EAR file to import. Use **Browse** to locate the EAR file and specify its full path name.
  - c. Optional: Specify a new enterprise application project name. A project name is assigned automatically. The project name you specify must be unique within the directory.
  - d. Specify the project location. Use **Browse** to locate the directory for the project files.
  - e. **5.1+** Select a target server. To use the Application profiling or Last participant support application assembly services of WebSphere Business Integration (WBI) Server Foundation, select the **Integration Server** target server.
  - f. Click **Finish**.
4. Verify the contents of the new enterprise application project in either of the following ways:
  - In the **J2EE Hierarchy** view, expand **Enterprise Applications** and view the new project.
  - Click **Window > Show View > Navigator** to see the associated files for the enterprise application project in a Navigator view.

## Importing WAR files

You can import a Web application archive (WAR) file and define a new enterprise archive (EAR) project and Web module for the WAR file using the Assembly Toolkit.

1. Start the Assembly Toolkit.
2. Click **File > Import**. Alternatively, you can right-click **Web Modules** in a view such as the J2EE Hierarchy view and click **Import**. Or, on Windows platforms, you can drag the WAR file and drop it on a view.

3. In the Import dialog, specify a WAR file and a Web project name:
  - a. Click **WAR file > Next**.
  - b. Specify a WAR file. Use **Browse** to locate the WAR file and specify its full path name.
  - c. Specify a Web project name. For example, if you are importing the HelloWorld.war file, you might name the project HelloWorld. Click **New** and specify HelloWorld for the project name.
  - d. Optional: To add the WAR file and Web project to an enterprise application, enable **Configure advanced options** and click **Next**. On the J2EE Settings page, specify the EAR project, the context root (Web project), and the J2EE 1.2 or 1.3 specification to use for the Web module. The J2EE 1.3 specification (the default) includes the Servlet 2.3 specification and the JSP 1.2 specification; applications developed for the J2EE 1.3 specification typically target a WebSphere Application Server Version 5.x server. Then, click **Next**. On the Features page, specify a feature for the Web project. For example, enable **Default synchronization policy for CVS repository** to have a .cvsignore file generated for the WEB-INF/classes directory. Then, click **Next** or **Finish**.
  - e. Click **Finish**.
4. Verify the contents of the new Web module in either of the following ways:
  - In the J2EE Hierarchy view, expand **Web Modules** and view the new module.
  - Click **Window > Show View > Navigator** to see the associated files for the Web module in a Navigator view.

## Importing client applications

You can import an application client JAR file into a new or existing enterprise application using the Assembly Toolkit.

1. Start the Assembly Toolkit.
2. Click **File > Import**. Alternatively, you can right-click **Application Client Modules** in a view such as the J2EE Hierarchy view and click **Import**. Or, on Windows platforms, you can drag the application client JAR file and drop it on a view.
3. In the Import dialog, specify the application client file and project name:
  - a. Click **App Client JAR file > Next**.
  - b. Specify the application client JAR file to be imported. Use **Browse** to locate the JAR file and specify its full path name.
  - c. Specify an application client project name. For example, if you are importing the HelloWorld.jar file, you might name the project HelloWorld. Click **New** and specify HelloWorld for the project name.
  - d. **5.1+** Select a target server. To use the Application profiling or Last participant support application assembly services of WebSphere Business Integration (WBI) Server Foundation, select the **Integration Server** target server.
  - e. Specify the enterprise archive (EAR) file into which to import the application client project.
  - f. Click **Finish**.
4. Verify the contents of the new application client module in either of the following ways:
  - In the J2EE Hierarchy view, expand **Application Client Modules** and view the new module.
  - Click **Window > Show View > Navigator** to see the associated files for the application client module in a Navigator view.

## Importing EJB files

You can import an enterprise bean (EJB) JAR file and define a new enterprise archive (EAR) project and EJB module for the enterprise bean JAR file using the Assembly Toolkit.

1. Start the Assembly Toolkit.

2. Click **File > Import**. Alternatively, you can right-click **EJB Modules** in a view such as the J2EE Hierarchy view and click **Import**. Or, on Windows platforms, you can drag the enterprise bean JAR file and drop it on a view.
3. In the Import dialog, specify the EJB JAR file and project name:
  - a. Click **EJB JAR file > Next**.
  - b. Specify the enterprise bean JAR file to import. Use **Browse** to locate the JAR file and specify its full path name.
  - c. **5.1+** Select a target server. To use an application assembly service of WebSphere Business Integration (WBI) Server Foundation, select the **Integration Server** target server. Available WBI assembly services include ActivitySession, Application profiling, Extended access, Extended messaging, Internationalization, Last participant support, and Pushdown.
  - d. Specify an EJB project name. For example, if you are importing the HelloWorld.jar file, you might name the project HelloWorld. Click **New**, specify HelloWorld for the project name, specify whether you want to use the EJB 1.1 or 2.0 specification (EJB 2.0 is the default), and click **Next**.
  - e. Name the enterprise archive (EAR) file into which to import the enterprise bean JAR file. The name must be unique among EAR files in the directory.
  - f. Click **Finish**.
4. Verify the contents of the new EAR file and EJB module in either of the following ways:
  - In the J2EE Hierarchy view, expand **Enterprise Applications** or **EJB Modules** and view the new modules.
  - Click **Window > Show View > Navigator** to see the associated files for the EAR file and EJB module in a Navigator view.

## Importing RAR files or connectors

You can import a resource adapter archive (RAR) file, or connector, and define a new enterprise archive (EAR) project and connector module using the Assembly Toolkit.

1. Start the Assembly Toolkit.
2. Right-click **Connector Modules** in a view such as the J2EE Hierarchy view and click **Import > Import Connector Module**.
3. In the Import dialog, specify the connector file and project name:
  - a. Specify the name of the RAR file to import. Use **Browse** to locate the RAR file and specify its full path name.
  - b. Specify a connector project name. For example, if you are importing the HelloWorld.rar file, you might name the project HelloWorld. Click **New**, specify HelloWorld for the project name, and click **Next**.
  - c. **5.1+** Select a target server. To use the Application profiling or Last participant support application assembly services of WebSphere Business Integration (WBI) Server Foundation, select the **Integration Server** target server.
  - d. If you want the connector project not to be part of an enterprise application, specify **Standalone connector project**.
  - e. If the connector project is not to stand alone, name the enterprise archive (EAR) file into which to import the RAR file. The name must be unique among EAR files in the directory.
  - f. If you want the Assembly Toolkit to overwrite existing resource files without first warning you that the files are changing, specify **Overwrite existing resources without warning**. The default is not to overwrite files without warning.
  - g. Click **Finish**.
4. Verify the contents of the new connector module in either of the following ways:



- In the J2EE Hierarchy view, expand **Enterprise Applications** or **Connector Modules** and view the new modules.
- Click **Window > Show View > Navigator** to see the associated files for the connector module in a Navigator view.

---

## Creating enterprise applications

Before you can deploy your archive files onto an application server, you must assemble them in an enterprise application archive (EAR) file. This article describes how to create a Java 2 Platform, Enterprise Edition (J2EE) enterprise application project using the Assembly Toolkit. After you create an enterprise application project, you can add (import) archive files such as Web application archives (WAR), resource adapter archives (RAR), enterprise bean (EJB) JAR files, and application client archives (JAR) files.

1. Start the Assembly Toolkit.
2. Click **File > New > Project**. Or, if you are working in the J2EE perspective, click **File > New > Enterprise Application Project** and skip step 3a below.
3. In the New Project dialog, create an enterprise application project:
  - a. Click **J2EE > Enterprise Application Project > Next**.
  - b. Specify whether you want an EAR file that supports J2EE 1.2 or 1.3, and click **Next**.
  - c. On the Enterprise Application Project page:
    - 1) Specify an EAR file name and location.
    - 2) To change the default project location, click **Browse** and specify a new location.
    - 3) **5.1+** Select a target server.

To use an application assembly service of WebSphere Business Integration (WBI) Server Foundation, select the **Integration Server** target server. Available WBI assembly services include ActivitySession, Application profiling, Extended access, Extended messaging, Internationalization, Last participant support, and Pushdown.
    - 4) Click **Next**.
  - d. Optional: On the EAR Module Projects page, select the existing modules that you want to add to the new enterprise application project. To create new modules for this enterprise application, click **New Module**. On the New Module Project page, select **Create default module projects** to create modules for application client, enterprise bean (EJB), Web or connector projects. You can use the default project names for the modules or specify different project names. If you clear the **Create default module projects** check box, you can select a single module type and proceed with the proper wizard for that project type. Then, click **Finish** to create the project modules and add their names to the list of available modules on the EAR Module Projects page.
  - e. Click **Finish**.
  - f. Optional: Confirm that you want to view the J2EE Hierarchy view.
4. Verify the contents of the new enterprise application in either of the following ways:
  - In the **J2EE Hierarchy** view, expand **Enterprise Application** and view the new EAR file.
  - Click **Window > Show View > Navigator** to see the associated files for the enterprise application in a Navigator view.

---

## Creating Web applications

In the Assembly Toolkit, you create and maintain resources for Web applications in Web projects. There are two types of Web projects, dynamic and static. Dynamic web projects can contain dynamic J2EE resources such as servlets, JavaServer Pages (JSP) files, filters, and associated metadata, in addition to static resources such as images and HTML files. Static Web projects only contain static resources.

Dynamic Web projects are always imbedded in enterprise application projects. Creating a Web project in the Assembly Toolkit requires that an enterprise application (EAR) project exist, or the Assembly Toolkit



creates one for you. Creating a Web project updates the `application.xml` deployment descriptor of the specified enterprise application project to define the Web project as a module element. If you are importing a WAR file rather than creating a Web project new, the WAR Import wizard requires that you specify a Web project, which already requires an EAR project.

This article describes how to create a dynamic Web project using the Assembly Toolkit. For instructions on how to create a static Web project, see the Assembly Toolkit online help. In the Assembly Toolkit, click **Help > Help Contents > Assembly Toolkit information > Web development > Tasks > Working with Web projects > Creating new static Web projects**.

1. Start the Assembly Toolkit.
2. Optional: View an animation file that shows how to create a dynamic Web project using the Assembly Toolkit. In the Assembly Toolkit, click **Help > Help Contents > Assembly Toolkit information > Web development > Tasks > Working with Web projects > Creating new dynamic Web projects > Show Me**.
3. Open the J2EE perspective to work with J2EE projects. Click **Window > Open Perspective > Other > J2EE**.
4. Click **File > New > Dynamic Web Project**.
5. On the Dynamic Web Project page of the New Web Project dialog:
  - a. Specify a Web project (WAR file) name.
  - b. Specify a location for the WAR file. To change the default WAR files location, click **Browse** and specify a new location.
  - c. Decide whether you want to accept the defaults associated with a dynamic Web project or configure advanced options. If you want to accept the defaults, deselect **Configure advanced options**. Otherwise, select **Configure advanced options** and **Next**. Step 6 describes the defaults and advanced options for a dynamic Web project.
6. If you selected **Configure advanced options**, you can customize the Web project options:
  - a. Specify a new or existing enterprise application (EAR) project to be associated with your new Web project for purposes of deployment. If you want to add a Web project as a module to another enterprise application project in the future, open the `application.xml` editor for the enterprise application project and select **Add** on the General page.
  - b. Provide a **Context root** value. The context root is the Web application root, the top-level directory of your application when it is deployed to a Web server. The default value is the name of your Web project. You can change the context root after you create a project using the project Properties dialog, which you access from the project's context menu.
  - c. From the J2EE Level drop-down list, select the appropriate Sun Microsystems Servlet and JSP specification level for the dynamic elements you plan to include in your Web project. Any new servlets and JSP files that you expect to create should adhere to the latest specification level available; previous specification levels are offered to accommodate any legacy dynamic elements that you expect to import into the project.
  - d. Click **Next**.
  - e. Optional: On the Features Page, select one or more of the Web project features and click **Next**.
  - f. Optional: Select **Use a default Page Template for the Web Site** if you want your entire Web site to share a common page template. If you want to use one of the sample templates provided, select **Sample Template** and then choose one of the templates shown in the **Thumbnail** box. If you want to use a template of your own, select **User-defined Template** and then click **Browse** to select the template from the file system. The default is not to use a page template.
7. Click **Finish**. A new Web project is created, reflecting the J2EE folder structure that specifies the location of web content files, class files, class paths, the deployment descriptor, and supporting metadata.
8. Verify the contents of the new Web project in either of the following ways:

- In the **J2EE Hierarchy** view, expand **Enterprise Applications** and the enterprise application associated with your Web project to view the new WAR file.
- Click **Window > Show View > Navigator** to see the associated files for the Web project in a Navigator view.

You can now begin creating or importing content for your Web project using the New File wizards or the Import wizards available from the **File** menu.

---

## Creating application clients

Application client projects contain programs that run on networked client systems. An application client project is deployed as a JAR file.

In the Assembly Toolkit, you can create and add an application client project to a new or existing enterprise application project.

1. Start the Assembly Toolkit.
2. Open the J2EE perspective to work with J2EE projects. Click **Window > Open Perspective > Other > J2EE**.
3. Click **File > New > Application Client Project**.
4. In the Application Client project creation dialog:
  - a. Select the Java 2 Platform, Enterprise Edition (J2EE) specification version to which you want your project to adhere, and click **Next**.
  - b. Name the application client project and specify its location. To change the default project location, click **Browse** and specify a new location. If you specify a non-default project location that is already being used by another project, the project creation will fail.
  - c. **5.1+** Select a target server. To use application assembly services of WebSphere Business Integration (WBI) Server Foundation such as Application profiling, select the **Integration Server** target server.
  - d. Specify a new or existing enterprise application (EAR) project to be associated with your new application client project for purposes of deployment. Select an existing enterprise application project from the drop-down list or type a new project name. Or, click **New** and create a new enterprise application. Note that if you type a new EAR project name, the EAR project will be created in the default location with the lowest compatible J2EE version based on the version of the project being created. If you want to specify a different version or a different location for the enterprise application, you must click **New** and create a new enterprise application.
  - e. Optional: If you are creating a new enterprise application project or if you have no module dependencies to specify, skip this step. Otherwise, click **Next** to specify module and JAR file dependencies. On the Module Dependencies page, select dependent JAR files or modules within the associated enterprise application project. This updates the runtime class-path and Java project build path with the appropriate JAR files. Application client modules, EJB modules, and Web modules can all have dependencies on EJB modules or utility JAR files. Modules cannot depend on WAR or application client JAR files.
  - f. Click **Finish** to create the application client project.
5. Verify the contents of the new application client project in either of the following ways:
  - In the **J2EE Hierarchy** view, expand **Enterprise Applications** and the enterprise application associated with your application client project to view the new JAR file.
  - Click **Window > Show View > Navigator** to see the associated files for the application client project in a Navigator view.

After creating an application client project, you can edit the application client deployment descriptor if default properties are not sufficient. In the Client Deployment Descriptor editor, you can add enterprise bean, resource, or resource environment references as well as view and edit source code.

For detailed instructions on adding enterprise bean, resource, or resource environment references, see the Assembly Toolkit online help. In the Assembly Toolkit, click **Help > Help Contents > Assembly Toolkit information > J2EE application development > Tasks > Configuring application client modules with the client deployment descriptor editor**. Similar information is in the information center for IBM WebSphere Studio Application Developer at <http://publib.boulder.ibm.com/infocenter/wsphelp/index.jsp>. Click **WebSphere Studio Application Developer > J2EE development > Tasks > Configuring application client modules with the client deployment descriptor editor**.

---

## Creating EJB modules

In the Assembly Toolkit, you can create and test enterprise beans that conform to the distributed component architecture defined in the Sun Microsystems Enterprise JavaBeans (EJB) specification and that support extended functionality for WebSphere Application Server.

You can create enterprise beans (either with or without inheritance) such as session beans, container-managed persistence (CMP) entity beans, bean-managed persistence (BMP) entity beans, or message-driven beans. Using the EJB deployment descriptor editor of the Assembly Toolkit, you can set deployment descriptor and assembly properties for enterprise beans.

This article describes how to create an EJB project (or EJB module) using the Assembly Toolkit.

1. Start the Assembly Toolkit.
2. Open the J2EE perspective to work with J2EE projects. Click **Window > Open Perspective > Other > J2EE**.
3. Click **File > New > EJB Project**.
4. In the New EJB Project dialog:
  - a. Select the EJB specification version to which you want your EJB project to adhere, and click **Next**. If you plan on using EJB 2.0 enterprise beans, you must specify an EJB 2.0 project. You can add EJB 1.1 enterprise beans to an EJB 2.0 project. An EJB 2.0 project must exist in a J2EE 1.3 enterprise application project. Your available options can differ, depending on the J2EE preferences defined.
  - b. Name the EJB project and specify its location. To change the default project location, click **Browse** and specify a new location. If you specify a non-default project location that is already being used by another project, the project creation will fail.
  - c. **5.1+** Select a target server. To use an application assembly service of WebSphere Business Integration (WBI) Server Foundation, select the **Integration Server** target server. Available WBI assembly services include ActivitySession, Application profiling, Extended access, Extended messaging, Internationalization, Last participant support, and Pushdown.
  - d. Specify a new or existing enterprise application (EAR) project to be associated with your new EJB project for purposes of deployment. Select an existing enterprise application project from the drop-down list or type a new project name. Or, click **New** and create a new enterprise application. Note that if you type a new EAR project name, the EAR project will be created in the default location with the lowest compatible J2EE version based on the version of the project being created. If you want to specify a different version or a different location for the enterprise application, you must click **New** and create a new enterprise application.
  - e. Optional: If you are creating a new enterprise application project or if you have no module dependencies to specify, skip this step. Otherwise, click **Next** to specify module and JAR file

dependencies. On the Module Dependencies page, select dependent JAR files or modules within the associated enterprise application project. Note that this page is available only if you are using an existing enterprise application project.

- f. Click **Finish** to create the EJB project.
5. Verify the contents of the new EJB project in either of the following ways:
    - In the **J2EE Hierarchy** view, expand **Enterprise Applications** and the enterprise application associated with your EJB project to view the new JAR file.
    - Click **Window > Show View > Navigator** to see the associated files for the EJB project in a Navigator view.

After you have an EJB project to hold enterprise beans, you can do the following:

- Create or import enterprise beans to your EJB project.
- Add methods to the home and remote interfaces.
- Add custom finders.
- Add and define additional CMP fields.
- Add relationships.
- Edit the EJB deployment descriptor if default properties are not sufficient.
- Create EJB access beans and use them to create your client application.
- Map enterprise beans to RDB tables.

For detailed instructions on creating CMP fields or CMP finder methods for entity beans, relating CMP fields, adding methods to interfaces, or managing enterprise beans, see the Assembly Toolkit online help. In the Assembly Toolkit, click **Help > Help Contents > Assembly Toolkit information > Enterprise JavaBeans (EJB) development > Tasks**. Similar information is in the information center for IBM WebSphere Studio Application Developer at <http://publib.boulder.ibm.com/infocenter/wsphelp/index.jsp>. Click **WebSphere Studio Application Developer > J2EE development > Tasks > Developing EJB applications**.

---

## Creating connector modules

A *connector* is a J2EE component that provides access to Enterprise Information Systems (EIS), and must comply to the J2EE Connector architecture (JCA). An *Enterprise Information System (EIS)* is a set of related classes that lets an application access a resource such as data, or an application on a remote server, often called a resource adapter.

This article describes how to create a connector project using the Assembly Toolkit.

1. Start the Assembly Toolkit.
2. Open the J2EE perspective to work with J2EE projects. Click **Window > Open Perspective > Other > J2EE**.
3. Click **File > New > Connector Project**.
4. In the New Connector Project dialog:
  - a. If you want your connector project to be a stand-alone project, select **Standalone connector project**. Selecting that the connector project stand alone disables the enterprise archive (EAR) project option at the bottom of the page.
  - b. Name the connector project and specify its location. To change the default project location, click **Browse** and specify a new location.
  - c. **5.1+** Select a target server. To use the Application profiling or Last participant support application assembly services of WebSphere Business Integration (WBI) Server Foundation, select the **Integration Server** target server.
  - d. If you specified that the connector project stand alone, specify a new or existing enterprise application (EAR) project to be associated with your new connector project for purposes of deployment. Select an existing enterprise application project from the drop-down list or type a new

project name. Or, click **New** and create a new enterprise application. Note that if you type a new EAR project name, the EAR project will be created in the default location with the lowest compatible J2EE version based on the version of the project being created. If you want to specify a different version or a different location for the enterprise application, you must click **New** and create a new enterprise application.

- e. Optional: If you are creating a new enterprise application project or if you have no module dependencies to specify, skip this step. Otherwise, click **Next** to specify module and JAR file dependencies. On the Module Dependencies page, select dependent JAR files or modules within the associated enterprise application project.
  - f. Click **Finish** to create the connector project.
5. Verify the contents of the new connector project in either of the following ways:
- In the **J2EE Hierarchy** view, expand **Enterprise Applications** and the enterprise application associated with your connector project to view the new JAR file.
  - Click **Window > Show View > Navigator** to see the associated files for the connector project in a Navigator view.

---

## Editing deployment descriptors

A deployment descriptor is an extensible markup language (XML) file that describes how to deploy a module or application by specifying configuration and container options. When you create a module, the Assembly Toolkit creates deployment descriptor files for the module.

You can edit a deployment descriptor file manually. However, it is preferable to edit a deployment descriptor using an Assembly Toolkit deployment descriptor editor to ensure that the deployment descriptor has valid properties and that its references contain appropriate values.

Deployment descriptor editor	Resources modified in the editor
Application deployment descriptor editor	<ul style="list-style-type: none"> <li>• application.xml</li> <li>• ibm-application-bnd.xmi</li> <li>• ibm-application-ext.xmi</li> </ul>
Web deployment descriptor editor	<ul style="list-style-type: none"> <li>• WEB-INF/web.xml</li> <li>• Binding information</li> <li>• IBM binding and extensions information such as ibm-web-bnd.xmi and ibm-web-ext.xmi files</li> </ul>
Enterprise bean (EJB) deployment descriptor editor	<ul style="list-style-type: none"> <li>• ejb-jar.xml</li> <li>• ibm-ejb-jar-bnd.xml</li> <li>• ibm-ejb-jar-ext.xml</li> <li>• ibm-ejb-access-bean.xml</li> </ul>
Client deployment descriptor editor	<ul style="list-style-type: none"> <li>• application-client.xml</li> <li>• ibm-application-client-bnd.xmi</li> <li>• ibm-application-client-ext.xmi</li> </ul>
Web services editor	<ul style="list-style-type: none"> <li>• webservicex.xml</li> <li>• ibm-webservicex-bnd.xmi</li> <li>• ibm-webservicex-ext.xmi</li> </ul>
Web services client editor	<ul style="list-style-type: none"> <li>• webservicexclient.xml</li> <li>• ibm-webservicexclient-bnd.xmi</li> <li>• ibm-webservicexclient-ext.xmi</li> </ul>

1. Ensure that you are working in the J2EE Perspective. Click **Window > Open Perspective > J2EE**.
2. In a J2EE Hierarchy view (**Window > Show View > J2EE Hierarchy**), right-click the module with deployment descriptor values that you want to browse or edit, and click **Open With > Deployment Descriptor Editor**. A deployment descriptor editor for the module displays in a view. You can click tabs such as **Overview**, **Module**, **Security**, and **Source** at the bottom of the view to browse or edit specific



deployment descriptor values. Clicking **Source** displays editable source code; it is preferable to edit values in fields on or accessible from the other tabs rather than edit the source code manually.

**5.1+** If you selected the **Integration Server** target server for a module, tabs for application assembly services of WebSphere Business Integration (WBI) Server Foundation also display. Available WBI assembly services include ActivitySession, Application profiling, Extended access, Extended messaging, Internationalization, Last participant support, and Pushdown.

3. Edit the deployment descriptor values as desired. For information on fields in the deployment descriptor editor, press F1 and click a topic.
4. Save your changes to the deployment descriptor.
  - a. Close the deployment descriptor editor.
  - b. When prompted, click **Yes** to indicate that you want to save changes to the deployment descriptor.

You also can save changes to deployment descriptors at any time by pressing Ctrl-S.

---

## Mapping enterprise beans to database tables

You can map enterprise bean JAR files (EJB modules) to relational database (RDB) tables using the EJB to RDB Mapping wizard of the Assembly Toolkit. The wizard creates EJB to RDB mappings for the following situations:

### Existing enterprise bean but no database schema

*Top Down* mapping generates a default database schema and a mapping from one or more existing enterprise beans.

### Existing database schema but no enterprise bean

*Bottom Up* mapping generates one or more enterprise beans and mappings from an existing database schema.

### Existing enterprise bean and database schema

*Meet In the Middle* mapping matches existing enterprise beans with existing database tables. You can match by name, by name and type, or by neither.

1. In the J2EE Hierarchy view, right-click the EJB module.
2. Click **Generate > EJB to RDB Mapping**.
3. After the wizard opens, press F1 and select a type of mapping. The online help provides detailed information on generating a mapping.
4. For EJB 2.0 projects, on the EJB to RGB Mapping page specify whether you want to create a new backend (*Top Down*) or use an existing backend (*Bottom Up* or *Meet In the Middle*) where the schema exists in the backend but without a mapping file. If you previously generated a mapping, you can create and map unmapped elements or open the mapping editor to manually make changes. In EJB 2.0, your mapping and schema files make up a *backend* for EJB 2.0 projects. You can have multiple backend folders for each project; for example, one DB2 and one Oracle backend. The wizard uses one database backend only as the default, but you can define as many as you need.
5. Follow the instructions in the wizard and in the online help.
6. Click **Finish** to generate the mapping.

## Mapping constraints for databases

Each of your Java 2 platform, Enterprise Edition (J2EE) applications (in an EAR file) can have more than one enterprise bean Container Managed Persistence (CMP) module in it. Each EJB CMP module (as a JAR file) can contain more than one CMP bean. You can configure each bean to connect to its own data source by using the different Java Naming and Directory Interface (JNDI) names of each data source.

However, there are some constraints:

- The deployment tool only enables you to map EJB CMP beans to different databases at the EJB JAR level. There are different backend ids for different database schema mapping. The product uses the attribute *backendId* to distinguish database specific code. For example, the code that accesses an

Oracle database is different from that accessing a DB2 database. Even for the same vendor, different versions of their products can result in different generated code.

- At run time, WebSphere Application Server uses the *CurrentBackendId* attribute to access the correct back end related code. You can only configure the *CurrentBackendId* at the EJB JAR level. This means that **all** of the EJB CMP beans within a specific EJB JAR file (which contains all the CMR beans) must connect to the same type of data source (that is, the same vendor, same version).
- Even though you can connect EJB CMP beans to different databases of the same type, the process fails if your connections result in a *table join* between tables from different databases. Various actions such as EJB Query , Inheritance support, *readAhead..*, and so on, can result in a table join situation.
- The datasource needs to be an XA datasource to support multiple databases' connection within a transaction.

---

## Verifying archive files

The Assembly Toolkit validates code when you request code validation manually, automatically during a resource change, and automatically during a build.

As part of validating the code, the validation checks for the following:

- Required deployment properties contain values.
  - Values specified for environment entries match their associated Java types.
  - In both enterprise archive (EAR) and Web archive (WAR) files:
    - The target enterprise bean of the link exists for enterprise bean (EJB) references.
    - The target role exists for security role references.
    - Security roles are unique.
  - Each module listed in the deployment descriptor exists in the archive for EAR files.
  - Files for icons, servlets, error and welcome pages listed in the deployment descriptor have corresponding files in the archive for WAR files.
  - For EJB modules:
    - All class files referenced in the deployment descriptor exist in the JAR file.
    - Method signatures for enterprise bean home, remote and implementation classes are compliant with the EJB 2.0 specification.
1. Optional: Specify whether you want automatic code validation during a resource change or during a build. The default is for automatic code validation.
    - a. In the J2EE Hierarchy view, right-click on a project.
    - b. Click **Properties > Validation**.
    - c. Ensure that the **Run validation** options for builds and for automatic validation are selected. Select **Override validation preferences** to disable automatic code validation.
    - d. If you changed the **Validation** settings, click **Apply** or **OK**.
  2. Optional: Specify validation options for a project. The default is to check all validators for a project during code validation. For an enterprise application project, the validators might be for DTD, EAR, Web services, XML, XML schema, or XSL files.
    - a. In the J2EE Hierarchy view, right-click the project containing the code that you want to validate.
    - b. Click **Properties > Validation**.
    - c. Select **Override validation preferences**.
    - d. Select the validators you want checked during code validation.
    - e. If you changed the **Validation** settings, click **Apply** or **OK**.
  3. Right-click the project containing the code that you want to validate and click **Run Validation** to manually validate the code.

The results of the code validation are shown in a Tasks view. For information on the results, select an entry in the Tasks view, press F1, and click **Tasks view**.



If your application module contains EJB modules, generate code for deployment. Otherwise, you are ready to deploy the application module (or stand-alone Web module) onto the application server.

---

## Generating code for EJB deployment

This task assumes you have already assembled an enterprise bean (EJB) module, added it to an application, saved the application, and verified the application.

Before installing your application in WebSphere Application Server, you must generate deployment code for the application. This step is required for EJB modules and for any enterprise application archive (EAR) files that contain EJB modules. During code generation, the Assembly Toolkit prepares entity bean (JAR) files for deployment in a run-time environment. If your EJB project contains container-managed persistence (CMP) beans that have not been mapped, generating deployment code creates a default top-down mapping.

1. If you have turned automatic validation off, manually validate your enterprise beans before generating deployment code for them. If validating your beans results in compilation errors or validation errors, fix the errors before generating deployment code. However, if validating your beans results in warning or information messages, you can generate deployment code.
2. If you have changed the class path of your EJB project, ensure that the source folder for your EJB project is at the beginning of the class path of the project. Generating deployment code imports both the JAR file and the source code of the JAR file, so entries on the class path must be in the correct order.
3. In the J2EE Hierarchy view of the Assembly Toolkit, right-click on the enterprise application (EAR file) or EJB module (enterprise bean JAR file) for which you want to generate code for deployment.
4. Click a **Generate Deployment Code** option.
  - For EAR files, click **Generate Deployment Code**.
  - For enterprise bean JAR files, click **Generate > Deployment and RMIC Code > EJB\_module > Finish**.

Alternatively, you can generate deployment code for enterprise bean JAR files using the deployment tool for Enterprise JavaBeans (ejbdeploy) from a command prompt. For a detailed list of available options in the EJBDeploy tool, enter `ejbdeploy` from a command prompt.

Code is generated into the folder where your enterprise beans are located. Problems with the generation of Java RMI stub compiler (RMIC) code result in a window that displays error messages.

Install the Java 2 Platform, Enterprise Edition (J2EE) application on your server machine. You can install the application onto a server using the administrative console. Before installing the application, you might need to set class paths.

---

## Generating code for Web service deployment

This task assumes you have already assembled a module enabled with Web services, added it to an application, saved the application, and verified the application.

Before installing your application in WebSphere Application Server, you must generate deployment code for the application. This step is required for Web services-enabled modules and for any enterprise application archive (EAR) files that contain Web services-enabled modules.

1. If you have turned automatic validation off, manually validate any modules that use Web services with the JSR109 Web services validator before generating deployment code for them. If validating your module results in compilation errors or validation errors, fix the errors before generating deployment code. However, if validating your module results in warning or information messages, you can generate deployment code.

2. In the J2EE Hierarchy view of the Assembly Toolkit, right-click on the Web services-enabled module (WAR, enterprise bean JAR, or application client JAR file) for which you want to generate code for deployment.
3. Click **Web Services > Deploy Web Service**. Alternatively, you can generate deployment code for Web services-enabled modules using the deployment tool for Web services (wsdeploy) from a command prompt.
4. If messages indicate that automatic file overwriting is not enabled, click **Yes to All** so the generated files are added to the module.
5. If errors such as *Unbound classpath variable: WAS\_50\_PLUGINDIR* appear in the Tasks list, change the Java build path libraries properties to define that variable to be the WebSphere Application Server installation directory.

Install the Java 2 Platform, Enterprise Edition (J2EE) application on your server machine. You can install the application onto a server using the administrative console. Before installing the application, you might need to set class paths.

---

## Assembly Toolkit: Resources for learning

Use the following links to find relevant supplemental information about the Assembly Toolkit. The information resides on IBM and non-IBM Internet sites, whose sponsors control the technical accuracy of the information.

These links are provided for convenience. Often, the information is not specific to the IBM WebSphere Application Server product, but is useful all or in part for understanding the product. When possible, links are provided to technical papers and Redbooks that supplement the broad coverage of the release documentation with in-depth examinations of particular product areas.

View links to additional information about:

- Programming instructions and examples
- Programming specifications
- Administration

### Programming instructions and examples

- Developing and testing a complete J2EE "Hello World" application with WebSphere Studio V5
- Getting to know WebSphere Studio Application Developer: Its capabilities, technologies, and relationship to the open-source Eclipse IDE
- Developing and Deploying an End-to-end J2EE Application to JBoss Application Server using WebSphere Studio V5
- JMS Applications with WebSphere Studio V5 -- Part 1: Developing a JMS Point-to-Point Application
- WebSphere Studio Version 5 Tips and Techniques
- Java 2 Enterprise Edition: Books index

### Programming specifications

- J2EE 1.3 specification
- EJB specifications
- Servlet specifications

### Administration

- Application Client files
- Connector RAR files



---

## Chapter 30. Deploying and managing applications

After you develop an enterprise application and configure an application server, you can use the administrative console to install application files on the server and manage the activity of deployed applications.

1. Install your application on your application server.
2. Start and stop applications.
3. Edit the administrative configuration for an application. Go to the settings page for an application, change the values for settings as needed, and click **OK**.
4. Export applications.
5. Export DDL files.
6. Update application binary files.
7. Uninstall applications.

After making changes to administrative configurations of your applications, ensure that you click **Save** on the administrative console taskbar to save the changes.

---

### Enterprise applications

Enterprise applications (or J2EE applications) are applications that conform to the Java 2 Platform, Enterprise Edition, specification.

Enterprise applications can consist of the following:

- Zero or more EJB modules
- Zero or more Web modules
- Zero or more connector modules (packaged in RAR files)
- Zero or more application client modules
- Additional JAR files containing dependent classes or other components required by the application
- Any combination of the above

A J2EE application is represented by, and packaged in, an enterprise archive (EAR) file.

---

### Installing a new application

To install an enterprise application to a WebSphere Application Server configuration, you can use the administrative console or the wsadmin tool. The steps below describe how to use the administrative console to install an application, EJB component, or Web module.

**Note:** Once you start performing the steps below, click **Cancel** to exit if you decide not to install the application. Do not simply move to another administrative console page without first clicking **Cancel** on an application installation page.

1. Click **Applications > Install New Application** in the console navigation tree. The first of two Preparing for application install pages is shown.
2. On the first Preparing for application install page:
  - a. Specify the full path name of the source application file (.ear file otherwise known as an EAR file). The EAR file that you are installing can be either on the client machine (the machine that runs the Web browser) or on the server machine (the machine to which the client is connected). If you specify an EAR file on the client machine, then the administrative console uploads the EAR file to the machine on which the console is running and proceeds with application installation. You can also specify a stand-alone WAR or JAR file for installation.
  - b. If you are installing a stand-alone WAR file, specify the context root.

- c. Click **Next**.
3. On the second Preparing for application install page:
  - a. Select whether to generate default bindings. Using the default bindings causes any incomplete bindings in the application to be filled in with default values. Existing bindings are not altered. You can customize default values used in generating default bindings. For example, you can specify JNDI prefix for all the EJB files in EJB modules, default data source and connection factory settings for EJB modules, virtual host for web modules, and so on. "Preparing for application installation settings" on page 1413" describes available customizations and provides sample bindings.
  - b. Click **Next**. The Install New Application pages are now shown. If you chose to generate default bindings, you can proceed to the Summary step (step 23 below). "Example: Installing an EAR file using the default bindings" on page 1417" provides sample steps.
4. On the **Step: Provide options to perform the installation** panel, provide values for the following settings specific to WebSphere Application Server. Default values are used if you do not specify a value.
  - a. For **Pre-compile JSP**, specify whether to precompile JSP files as a part of installation. The default is not to precompile JSP files.
  - b. For **Directory to Install Application**, specify the directory to which the application EAR file will be installed. The default value is the value of APP\_INSTALL\_ROOT/*cell\_name*, where the APP\_INSTALL\_ROOT variable is *install\_root/installedApps*; for example, C:\WebSphere\AppServer\installedApps\*cell\_name*.
 

**Note:** If an installation directory is not specified when an application is installed on a single-server (base) configuration, the application is installed in APP\_INSTALL\_ROOT/*base\_cell\_name*. When the base server is made a part of a Network Deployment configuration (using the addNode utility), the cell name of the new configuration becomes the cell name of the deployment manager node. If the -includeapps option is used for the addNode utility, then the applications that are installed prior to the addNode operation still use the installation directory APP\_INSTALL\_ROOT/*base\_cell\_name*. However, an application that is installed after the base server is added to the network configuration uses the default installation directory APP\_INSTALL\_ROOT/*network\_cell\_name*. To move the application to the APP\_INSTALL\_ROOT/*network\_cell\_name* location upon running the addNode operation, you should explicitly specify the installation directory as \${APP\_INSTALL\_ROOT}/\${CELL} during installation. In such a case, the application files can always be found under APP\_INSTALL\_ROOT/*current\_cell\_name*.
  - c. For **Distribute Application**, specify whether WebSphere Application Server expands or deletes application binaries in the installation destination. The default is to enable application distribution. As a result, when you save changes in the console, application binaries for newly installed applications are expanded to the directory specified. The binaries are also deleted when you uninstall and save changes to the configuration. If you disable this option, then you must ensure that the application binaries are expanded appropriately in the destination directories of all nodes where the application is expected to run.
  - d. For **Use Binary Configuration**, specify whether the application server uses the binding, extensions, and deployment descriptors located with the application deployment document, the deployment.xml file (default), or those located in the EAR file. The default is not to use the binary configuration.
  - e. For **Deploy EJBs**, specify whether the EJBDeploy tool runs during application installation. The tool generates code needed to run EJB files. The default is not to run the EJBDeploy tool. You must enable this setting if the EAR file was assembled using the Assembly Toolkit and the EJBDeploy tool was not run during assembly, if the EAR file was not assembled using the Assembly Toolkit tool, or if the EAR file was assembled using versions of the Application Assembly Tool (AAT) previous to Version 5. Note that enabling this setting might cause the installation program to run for several minutes.

- f. For **Application Name**, name the application. Application names must be unique within a cell and cannot contain characters that are not allowed in object names (see "Object names" in the Information Center).
  - g. For **Create MBeans for Resources**, specify whether to create MBeans for various resources (such as servlets or JSP files) within an application when the application is started. The default is to create MBean instances.
  - h. For **Enable class reloading**, specify whether to enable class reloading when application files are updated. The default is not to enable class reloading. For EJB modules or any non-Web modules, enabling class reloading sets `reloadEnabled` to `true` in the `deployment.xml` file for the application. If an application's class definition changes, the application server run time stops and starts the application to reload application classes.  
 For Web modules such as servlets and JavaServer page (JSP) files, a Web container reloads a Web module only when the IBM extension `reloadingEnabled` in the `ibm-web-ext.xmi` file is set to `true`. You can set `reloadingEnabled` to `true` when editing your Web module's extended deployment descriptors in an assembly tool.  
 To disable reloading of a Web module, set the IBM extension `reloadingEnabled` in the `ibm-web-ext.xmi` file to `false`. Or, if the Web module has the IBM extension `reloadingEnabled` in the `ibm-web-ext.xmi` file set to `true`, enable class loading, and set the **Reload Interval** property to zero (0).
  - i. For **Reload Interval**, specify the number of seconds to scan the application's file system for updated files. The default is the value of the reload interval attribute in the IBM extension (META-INF/ibm-application-ext.xmi) file of the EAR file. To enable reloading, specify a value greater than zero (for example, 1 to 2147483647). To disable reloading, specify zero (0).  
 The reload interval specified here overrides the value specified in the IBM extensions for each Web module in the EAR file (which in turn overrides the reload interval specified in the IBM extensions for the application in the EAR file). This setting takes effect only if class reloading is enabled.
  - j. For **Deploy WebServices**, specify whether the Web services deploy tool `wsdeploy` runs during application installation. The tool generates code needed to run applications using Web services. The default is not to run the `wsdeploy` tool. You must enable this setting if the EAR file contains modules using Web services and has not previously had the `wsdeploy` tool run on it, either from the **Web Services > Deploy Web Service** menu of the Assembly Toolkit or from a command line.
5. If your application uses EJB modules, on the **Step: Provide JNDI Names for Beans** panel, specify a JNDI name for each enterprise bean in every EJB module. You must specify a JNDI name for every enterprise bean defined in the application. For example, for the EJB module `MyBean.jar`, specify `MyBean`.
  6. If your application uses EJB modules that contain Container Managed Persistence (CMP) beans that are based on the EJB 1.x specification, for **Step: Provide default datasource mapping for modules containing 1.x entity beans**, specify a JNDI name for the default data source for the EJB modules. The default data source for the EJB modules is optional if data sources are specified for individual CMP beans.
  7. If your application has CMP beans that are based on the EJB 1.x specification, for **Step: Map datasources for all 1.x CMP**, specify a JNDI name for data sources to be used for each of the 1.x CMP beans. The data source attribute is optional for individual CMP beans if a default data source is specified for the EJB module that contains CMP beans. If neither a default data source for the EJB module nor a data source for individual CMP beans are specified, then a validation error displays after you click **Finish** (step 23) and the installation is cancelled.
  8. If your application defines EJB references, for **Step: Map EJB references to beans**, specify JNDI names for enterprise beans that represent the logical names specified in EJB references. Each EJB reference defined in the application must be bound to an EJB file before clicking **Finish** on the Summary panel.
  9. If your application defines resource references, for **Step: Map resource references to resources**, specify JNDI names for the resources that represent the logical names defined in resource



- references. Each resource reference defined in the application must be bound to a resource defined in your WebSphere Application Server configuration before clicking on **Finish** on the Summary panel.
10. If your application uses Web modules, for **Step: Map virtual hosts for web modules**, select a virtual host from the list that should map to a Web module defined in the application. The port number specified in the virtual host definition is used in the URL that is used to access artifacts such as servlets and JSP files in the Web module. Each Web module must have a virtual host to which it maps. Not specifying all needed virtual hosts will result in a validation error displaying after you click **Finish** on the Summary panel.
  11. On the **Step: Map modules to application servers** panel, for every module select a target server or a cluster from the **Clusters and Servers** list. Select the check box beside **Module** to select all of the application modules or select individual modules.
  12. If the application has security roles defined in its deployment descriptor then, for **Step: Map security roles to users/groups**, specify users and groups that are mapped to each of the security roles. Select the check box beside **Role** to select all of the roles or select individual roles. For each role, you can specify if predefined users such as **Everyone** or **All Authenticated users** are mapped to it. To select specific users or groups from the user registry:
    - a. Select a role and click **Lookup users** or **Lookup groups**.
    - b. On the Lookup users/groups panel shown, enter search criteria to extract a list of users or groups from the user registry.
    - c. Select individual users or groups from the results displayed.
    - d. Click **OK** to map the selected users or groups to the role selected on the **Step: Map security roles to users/groups** panel.
  13. If the application has Run As roles defined in its deployment descriptor, for **Step: Map RunAs roles to user**, specify the Run As user name and password for every Run As role. Run As roles are used by enterprise beans that must run as a particular role while interacting with another enterprise bean. Select the check box beside **Role** to select all of the roles or select individual roles. After selecting a role, enter values for the user name, password, and verify password and click **Apply**.
  14. If your application contains EJB 1.x CMP beans that do not have method permissions defined for some of the EJB methods, for **Step: Ensure all unprotected 1.x methods have the correct level of protection**, specify if you want to leave such methods unprotected or assign protection with deny all access.
  15. If your application contains message driven enterprise beans, for **Step: Provide Listener Ports for messaging beans**, provide a listener port name for every message driven bean. If a name is not specified for each bean, then a validation error displays after you click on **Finish** on the Summary panel.
  16. If your application uses EJB modules that contain CMP beans that are based on the EJB 2.0 specification, for **Step: Provide default datasource mapping for modules containing 2.0 entity beans**, specify a JNDI name for the default data source and the type of resource authorization to be used for the default data source for the EJB modules. The default data source for EJB modules is optional if data sources are specified for individual CMP beans.
  17. If your application has CMP beans that are based on the EJB 2.0 specification, on the **Step: Map datasources for all 2.0 CMP** panel, for each of the 2.0 CMP beans specify a JNDI name and the type of resource authorization for data sources to be used. The data source attribute is optional for individual CMP beans if a default data source is specified for the EJB module that contains CMP beans. If neither a default data source for the EJB module nor a data source for individual CMP beans are specified, then a validation error is shown after you click **Finish** and installation is cancelled.
  18. If your application contains EJB 2.0 CMP beans that do not have method permissions defined in the deployment descriptors for some of the EJB methods, on the **Step: Ensure all unprotected 2.0 methods have the correct level of protection** panel, specify whether you want to assign a specific role to the unprotected methods, add the methods to the exclude list, or mark them as unchecked. Methods added to the exclude list are marked as uncallable. For methods marked unchecked no authorization check is performed prior to their invocation.



19. If the **Deploy EJBs** setting is enabled on the **Provide options to perform the installation** panel, then you can specify options for the EJBDeploy tool on the **Step: Provide options to perform the EJB Deploy** panel. On this panel, you can specify extra class path, rmic options, database types, and database schema names to be used while running the EJBDeploy tool. The tool is run on the EAR file during installation after you click **Finish**.
20. If your application contains resource environment references, for **Step: Mapping Resource Environment References to Resources**, specify JNDI names of resources that map to the logical names defined in resource environment references. If each resource environment reference does not have a resource associated with it, a validation error is shown after you click **Finish**.
21. If your application defines **Run-As Identity** as *System Identity*, for **Step: Replacing RunAs System to RunAs Roles**, you can optionally change it to *Run-As role* and specify a user name and password for the Run As role specified. Selecting *System Identity* implies that the invocation is done using the WebSphere Application Server security server ID and should be used with caution as this ID has more privileges.
22. If your application has resource references that map to resources that have an Oracle database doing backend processing, for **Step: Specify the isolation level for Oracle type provider**, specify or correct the isolation level to be used for such resources when used by the application. Oracle databases support ReadCommitted and Serializable isolation levels only.
23. On the Summary panel, verify the cell, node, and server onto which the application modules will install. Beside the **Cell/Node/Server** option, click **Click here** and verify the settings. Then click **Finish**.  
**Note:** After clicking **Finish**, if you receive an OutOfMemory exception and the source application file does not install, your system might not have enough memory or your application might have too many modules in it to install successfully onto the server. If lack of system memory is not the cause of the exception, package your application again so the .ear file has fewer modules. If lack of system memory and the number of modules are not the cause of the exception, check the options you specified on the Java Virtual Machine page of the application server running the administrative console. Then, try installing the application file again.
24. Associate any shared libraries that the application needs to the application. (See "Managing shared libraries" in the Information Center.)
25. Click **Save** on the administrative console taskbar to save the changes to your configuration. The application is registered with the administrative configuration and application files are copied to the target directory, which is *install\_root/installedApps/cell\_name* by default or the directory that you designate. For the single-server (base) installation, application files are copied to the destination directory when you click **Save**; for the Network Deployment installation, files are copied to remote nodes when the configuration on the deployment manager synchronizes with the configuration on individual nodes.
26. Start the application.
27. Test the application. For example, point a Web browser at the URL for the deployed application and examine the performance of the application. If necessary, update the application.

## Preparing for application installation settings

Use this page to install an application (EAR file) or module (JAR or WAR file).

To view this administrative console page, click **Applications > Install New Application**.

Follow the steps on this page to install an application or module. You must complete, at minimum, the first step; you must complete some or all of the later steps, depending on whether you are installing an application, EJB module, or Web module.

### Path

Specifies the fully qualified path to the .ear, .jar, or .war file for the enterprise application.

Use **Local path** if the browser and application files are on the same machine (whether or not the server is on that machine, too).

Use **Server path** if the application file resides on any node in the current cell context. You can browse the entire file system of a node if the node agent or deployment manager is running on that selected node. Only .ear, .jar, or .war files are shown during the browsing.

During application installation, application files are typically uploaded from a client machine running the browser to the server machine running the administrative console, where they are deployed. In such cases, the Web browser running the administrative console is used to select EAR, WAR, or JAR modules to upload to the server machine.

In some cases, however, the application files reside on the file system of any of the nodes in a cell. To have the application server install these files, use the **Server path** option.

You can also use this option to specify an application file already residing on the machine running the application server. For example, the field value on Windows NT might be C:\WebSphere\AppServer\installableApps\test.ear. If you are installing a stand-alone WAR module, then you also must specify the context root.

### Context Root

Specifies the context root of the Web application (WAR).

This field is used only to install a stand-alone WAR file. The context root is combined with the defined servlet mapping (from the WAR file) to compose the full URL that users type to access the servlet. For example, if the context root is /gettingstarted and the servlet mapping is MySession, then the URL is http://host:port/gettingstarted/MySession.

### Generate Default Bindings

Specifies whether to generate default bindings. If you place a check mark in the check box, then any incomplete bindings in the application are filled in with default values. Existing bindings are not altered.

By choosing this option, you can directly jump to the Summary step and install the application if none of the steps have a red asterisk (\*) next to them. A red asterisk denotes that the step has incomplete data and requires a valid value. On the Summary panel, verify the cell, node and server on which the application is installed.

Bindings are generated as follows:

- EJB JNDI names are generated of the form *prefix/ejb-name*. The default prefix is *ejb*, but can be overridden. The *ejb-name* is as specified in the deployment descriptors `<ejb-name>` tag.
- EJB references are bound as follows: If an `<ejb-link>` is found, it is honored. Otherwise, if a unique enterprise bean is found with a matching home (or local home) interface as the referenced bean, the reference is resolved automatically.
- Resource reference bindings are derived from the `<res-ref-name>` tag. Note that this action assumes that the `java:comp/env` name is the same as the resource global JNDI name.
- Connection factory bindings (for EJB 2.0 JAR files) are generated based on the JNDI name and authorization information provided. This action results in default connection factory settings for each EJB 2.0 JAR file in the application being installed. No bean-level connection factory bindings are generated.
- Data source bindings (for EJB 1.1 JAR files) are generated based on the JNDI name, data source user name password options. This results in default data source settings for each EJB JAR file. No bean-level data source bindings are generated.
- Message-driven bean (MDB) listener ports are derived from the MDB `<ejb-name>` tag with the string `Port` appended.
- For .war files, the virtual host is set as `default_host` unless otherwise specified.

The default strategy suffices for most applications or at least for most bindings in most applications. However, it does not work if:

- You want to explicitly control the global JNDI names of one or more EJB files.
- You need tighter control of data source bindings for CMPs. That is, you have multiple data sources and need more than one global data source.
- You must map resource references to global resource JNDI names that are different from the `java:comp/env` name.

In such cases, you can alter the behavior with an XML document (a custom strategy). Use the **Specific bindings file** field to specify a custom strategy and see the field's help for examples.

## Prefixes

Specifies prefixes to use for generated JNDI names.

## Override

Specifies whether to override existing bindings.

If this check box is checked, the existing bindings are overridden by the generated ones.

## EJB 1.1 CMP bindings

Specifies the default data source JNDI name.

If the **Default Bindings for EJB 1.1 CMPs** radio button is selected, specify the JNDI name for the default data source to be used with the CMP 1.1 beans. Also specify the user ID and password for this default data source.

## Connection Factory Bindings

Specifies the default data source JNDI name.

If the **Default connection factory bindings** radio button is selected, specify the JNDI name for the default data source to be used with the bindings. Also specify the resource authorization.

## Virtual Host

Specifies the virtual host for WAR modules.

## Specific bindings file

Specifies a bindings file that overrides the default binding.

Alter the behavior of the default binding with an XML document (*aka* custom strategy). Custom strategies extend the default strategy so you only need to customize those areas where the default strategy is insufficient. That is, you only need to describe how you want to change the bindings generated by the default strategy; you do not have to define bindings for the entire application.

Brief examples of how to override various aspects of the default bindings generator follow:

### Controlling an EJB JNDI name

```
<?xml version="1.0"?>
<!DOCTYPE dfltbndngs SYSTEM "dfltbndngs.dtd">
<dfltbndngs>
 <module-bindings>
 <ejb-jar-binding>
 <jar-name>helloEjb.jar</jar-name>
<!-- this name must match the module name in the .ear file -->
 <ejb-bindings>
 <ejb-binding>
 <ejb-name>HelloEjb</ejb-name>
<!-- this must match the <ejb-name> entry in the EJB jar DD -->
 <jndi-name>com/acme/ejb/HelloHome</jndi-name>
```

```

 </ejb-binding>
 </ejb-bindings>
</ejb-jar-binding>
</module-bindings>
</dfldbndngs>

```

### Setting the connection factory binding for an EJB JAR file

```

<!DOCTYPE dfldbndngs SYSTEM "dfldbndngs.dtd">
<dfldbndngs>
 <module-bindings>
 <ejb-jar-binding>
 <jar-name>yourEjb20.jar</jar-name>
 <connection-factory>
 <jndi-name>eis/jdbc/YourData_CMP</jndi-name>
 <res-auth>Container</res-auth>
 </connection-factory>
 </ejb-jar-binding>
 </module-bindings>
</dfldbndngs>

```

### Setting the connection factory binding for an EJB file

```

<?xml version="1.0">
<!DOCTYPE dfldbndngs SYSTEM "dfldbndngs.dtd">
<dfldbndngs>
 <module-bindings>
 <ejb-jar-binding>
 <jar-name>yourEjb20.jar</jar-name>
 <ejb-bindings>
 <ejb-binding>
 <ejb-name>YourCmp20</ejb-name>
<!-- this matches the ejb-name tag in the DD -->
 <connection-factory>
 <jndi-name>eis/jdbc/YourData_CMP</jndi-name>
 <res-auth>PerConnFact</res-auth>
 </connection-factory>
 </ejb-binding>
 </ejb-bindings>
 </ejb-jar-binding>
 </module-bindings>
</dfldbndngs>

```

### Overriding a Resource Ref Binding from a WAR, EJB JAR file, or J2EE client JAR file

Example code for overriding a Resource Ref Binding from a WAR file follows. Use similar code to override a Resource Ref Binding from an enterprise bean (EJB) JAR file or a J2EE client JAR file.

```

<?xml version="1.0"?>
<!DOCTYPE dfldbndngs SYSTEM "dfldbndngs.dtd">
<dfldbndngs>
 <module-bindings>
 <war-binding>
 <jar-name>hello.war</jar-name>
 <resource-ref-bindings>
 <resource-ref-binding>
 <!-- the following must match the resource-ref in the DD -->
 <resource-ref-name>jdbc/MyDataSrc</resource-ref-name>
 <jndi-name>war/override/dataSource</jndi-name>
 </resource-ref-binding>
 </resource-ref-bindings>
 </war-binding>
 </module-bindings>
</dfldbndngs>

```

### Overriding MDB JMS listener ports

```

<?xml version="1.0"?>
<!DOCTYPE dfltbndngs SYSTEM "dfltbndngs.dtd">
<dfltbndngs>
 <module-bindings>
 <ejb-jar-binding>
 <jar-name>YourEjbJar.jar</jar-name>
 <ejb-bindings>
 <ejb-binding>
 <ejb-name>YourMDB</ejb-name>
 <listener-port>yourMdbListPort</listener-port>
 </ejb-binding>
 </ejb-bindings>
 </ejb-jar-binding>
 </module-bindings>
</dfltbndngs>

```

## Example: Installing an EAR file using the default bindings

An example of a simple .ear file installation using the default bindings follows:

1. Go to the Preparing for application install pages. Click **Applications > Install an Application** in the console navigation tree.
2. For **Path**, specify the full path name of the .ear file. For this example, the base file name is my\_appl.ear and the file resides on a server at C:\sample\_apps.
  - a. Select the **Server path** radio button and click **Browse**.
  - b. On the Browse Remote Filesystems page, click on the name of the node that holds the my\_appl.ear file, **C:\, sample\_apps, my\_appl.ear**, and then **OK**.
3. Now that a value is given for **Path**, on the first Preparing for application install page, click **Next**.
4. On the second Preparing for application install page, place a checkmark beside the **Generate Default Bindings** check box and click **Next**. Using the default bindings causes any incomplete bindings in the application to be filled in with default values. Existing bindings are not changed. By choosing this option, you can directly jump to the Summary step.
5. On the Install New Application page, click on **Summary**, the last step.
6. On the Summary panel, verify the cell, node, and server onto which the application files will install.
  - a. Beside the **Cell/Node/Server** option, click **Click here**.
  - b. On the **Map modules to application servers** panel, select the server onto which the application files will install from the **Clusters and Servers** list, place a checkmark in the check box beside **Module** to select all of the application modules, and click **Next**.

Because my\_appl.ear does not require any additional settings to complete an installation, the Summary panel displays again.
7. On the Summary panel, click **Finish**.

---

## Enterprise application collection

Use this page to view and manage enterprise applications.

To view this administrative console page, click **Applications > Enterprise Applications**.

### Name

Specifies the name of the installed (or deployed) application. Application names must be unique within a cell and cannot contain characters that are not allowed in object names.

### Status

Indicates whether the application deployed on the application server is started, stopped, or unavailable.

## Enterprise application settings

Use this page to configure an enterprise application.

To view this administrative console page, click **Applications > Enterprise Applications** > *application\_name*.

### Name

Specifies a logical name for the application. Application names must be unique within a cell and cannot contain characters that are not allowed in object names.

**Data type** String

### Starting Weight

Specifies the order in which applications are started when the server starts. The application with the lowest starting weight is started first.

**Data type** Integer  
**Default** 1  
**Range** 0 to 100

### Application Binaries

Specifies the directory to which the application EAR file will be installed. The default value is the value of `APP_INSTALL_ROOT/cell_name`, where the `APP_INSTALL_ROOT` variable is `install_root/installedApps`; for example, `C:\WebSphere\AppServer\installedApps\cell_name`.

You can specify an absolute path or use a pathmap variable such as `#{MY_APPS}`. You can use a pathmap variable in any installation though it is particularly needed when installing an application on a cluster with members on heterogeneous nodes because, in such cases, there might not be a single way to specify an absolute path. A WebSphere Application Server variable `#{CELL}` that denotes the current cell name can also be in the pathmap variable; for example, `#{MY_APP}/#{CELL}`.

You can define WebSphere Application Server variables on the WebSphere Variables page of the administrative console, accessed by clicking **Environment > Manage WebSphere Variables**.

**Data type** String  
**Units** Full path name

### Use Metadata From Binaries

Specifies whether the application server uses the binding, extensions, and deployment descriptors located with the application deployment document, the `deployment.xml` file (default), or those located in the enterprise application resource (EAR) file.

**Data type** Boolean  
**Default** false

### Enable Distribution

Specifies whether WebSphere Application Server expands or deletes application binaries in the installation destination. The default is to enable application distribution. Application binaries for installed applications are expanded to the directory specified. The binaries are also deleted when you uninstall and save changes to the configuration. If you disable this option, then you must ensure that the application binaries are expanded appropriately in the destination directories of all nodes where the application runs.

**Data type** Boolean  
**Default** true

## Classloader Mode

Specifies whether the class loader searches in the parent class loader or in the application class loader first to load a class. The standard for JDK class loaders and WebSphere Application Server class loaders is PARENT\_FIRST. By specifying PARENT\_LAST, your application can override classes contained in the parent class loader, but this action can potentially result in ClassCastException or LinkageErrors if you have mixed use of overridden classes and non-overridden classes.

The options are PARENT\_FIRST and PARENT\_LAST. The default is to search in the parent class loader before searching in the application class loader to load a class.

<b>Data type</b>	String
<b>Default</b>	PARENT_FIRST

## WAR Classloader Policy

Specifies whether to use a single class loader to load all WAR files of this application or to use a different class loader for each WAR file.

The options are APPLICATION and MODULE. The default is to use a separate class loader to load each WAR file.

<b>Data type</b>	String
<b>Default</b>	MODULE

## Create MBeans for Resources

Specifies whether to create MBean files for various resources (such as servlets or JSP files) within an application.

<b>Data type</b>	Boolean
<b>Default</b>	true

## Reload Enabled

Specifies whether to enable class reloading when application files are updated.

For EJB modules or any non-Web modules, setting **Reload Enabled** to true sets reloadEnabled to true in the deployment.xml file for the application. If an application's class definition changes, the application server run time stops and starts the application to reload application classes.

For Web modules such as servlets and JavaServer page (JSP) files, a Web container reloads a Web module only when the IBM extension reloadingEnabled in the ibm-web-ext.xmi file is set to true. You can set reloadingEnabled to true when editing your Web module's extended deployment descriptors in an assembly tool.

To enable reloading of a Web module, where you also want reloading of EJB and non-Web modules enabled:

1. Set the IBM extension reloadingEnabled in the ibm-web-ext.xmi file to true.
2. Set this **Reload Enabled** property to true.
3. Set the **Reload Interval** property to a value greater than zero (for example, 1 to 2147483647).

To enable reloading of a Web module only, and not enable reloading of EJB or non-Web modules:

1. Set the IBM extension reloadingEnabled in the ibm-web-ext.xmi file to true.
2. Set the IBM extension reload interval attribute in the ibm-web-ext.xmi file to a value greater than zero (for example, 1 to 2147483647).
3. Set this **Reload Enabled** property to false.



To disable reloading of a Web module, set the IBM extension reloadingEnabled in the `ibm-web-ext.xml` file to `false`. Or, if the Web module has the IBM extension reloadingEnabled in the `ibm-web-ext.xml` file set to `true`, to disable reloading using the administrative console:

1. Set this **Reload Enabled** property to `true`.
2. Set the **Reload Interval** property to zero (0).

<b>Data type</b>	Boolean
<b>Default</b>	false

## Reload Interval

Specifies the number of seconds to scan the application's file system for updated files. The default is the value of the reload interval attribute in the IBM extension (`META-INF/ibm-application-ext.xml`) file of the EAR file.

To enable reloading, specify a value greater than zero (for example, 1 to 2147483647). To disable reloading, specify zero (0).

The reload interval specified here overrides the value specified in the IBM extensions for each non-Web module in the EAR file (which in turn overrides the reload interval specified in the IBM extensions for the application in the EAR file). The reload interval attribute takes effect only if class reloading is enabled.

The range is from 0 to 2147483647.

<b>Data type</b>	Integer
<b>Units</b>	Seconds

## Target mapping collection

Use this page to manage mappings of deployed applications or modules to servers or clusters.

To view this administrative console page, click **Applications > Enterprise Applications > *application\_name* > Target Mappings**.

### **Target:**

States the name of the target server or cluster to which the application or module maps. You specify the target on the Map modules to application servers page accessed from the settings for an application.

### **Node:**

Specifies the node name if the target is a server.

### **Status:**

Indicates whether the status of the application running on the target server or cluster is started, stopped or unavailable.

### **Target mapping settings:**

Use this page to map a deployed application or module to a server or cluster.

To view this administrative console page, click **Applications > Enterprise Applications > *application\_name* > Target Mappings > *target\_name***.

### **Target:**

States the name of the target server or cluster to which the application or module maps. You specify the target on the Map modules to application servers page accessed from the settings for an application.

**Data type** String

*Enabled:*

Indicates whether the application modules installed on the target server are started (or enabled) when the server starts. This sets the initial state of application modules. A true value indicates that the corresponding modules are enabled and thus are accessible when the server starts. A false value indicates that the corresponding modules are not enabled and thus are not accessible when the server starts.

**Data type** Boolean  
**Default** true

---

## Starting and stopping applications

You can start an application that is not running (has a status of *Stopped*) or stop an application that is running (has a status of *Started*).

1. Go to the Enterprise Applications page. Click **Applications > Enterprise Applications** in the console navigation tree.
2. Check the check box for the application you want started or stopped.
3. Click a button:

Option	Description
<b>Start</b>	Runs the application and changes the state of the application from <i>Stopped</i> to <i>Started</i> .
<b>Stop</b>	Stops the processing of the application and changes the state of the application from <i>Started</i> to <i>Stopped</i> .

To restart a running application, place a check mark in the check box for the application you want to restart, click **Stop** and then click **Start**.

The status of the application changes and a message stating that the application started or stopped displays at the top the page.

---

## Exporting applications

You can export an enterprise application to a location of your choice. Exporting applications enables you to back up your applications and preserve binding information for the applications. You might export your applications before updating installed applications or migrating to a later version of the WebSphere Application Server product.

1. Click **Applications > Enterprise Applications** in the administrative console navigation tree to access the Enterprise Applications page.
2. Place a check mark in the check box beside the application and click **Export**.
3. On the Export Application EAR Files page, click on the link to download the exported EAR file.
4. Use the browser dialogue to specify a location at which to save the exported EAR file and click **OK**.

The file containing binding information is exported to the specified node and directory, and has the name *enterprise\_application\_name.ear*.

---

## Exporting DDL files

You can export the DDL files (Table.ddl) in the EJB modules of the application to a location of your choice.

1. Click **Applications > Enterprise Applications** in the administrative console navigation tree to access the Enterprise Applications page.
2. Place a checkmark in the check box beside the application and click **Export DDL**. If the application has no DDL files in any of its EJB modules, then the message *No DDL files were found* displays at the top of the page. If the application has DDL files in its EJB modules, then a page displays listing DDL files in the format `appName.ear/_module.jar_Table.ddl`.
3. Click on a file in the list to download the file to your machine.

---

## Updating applications

You can update an application deployed on a server. The steps below describe how to update a deployed application using the administrative console.

**Note:** You can also update applications using the wsadmin tool, which provides updating capabilities identical to those available using the administrative console. Further, in some situations, you can update applications without needing to restart the application server.

1. Update the contents of the application and reassemble it, using the Assembly Toolkit. Typical tasks include adding or editing assembly properties, adding or importing modules into an application, and adding enterprise beans, Web components, and files.
2. Go to the Applications page of the administrative console. Click **Applications > Enterprise Applications** in the console navigation tree.
3. Back up the application. Place a checkmark in the check box beside the application you want uninstalled and click **Export** to export the application to an EAR file and preserve the binding information.
4. With a checkmark beside the application, click **Update**. The binding information of the updated (new) version of the application merges with the binding information from the installed (old) version. Then, the older version uninstalls from the configuration and the new version installs.
5. Complete the steps in the Preparing for application update page and the pages that follow it. See information on installing applications and on the settings page for application installation for guidance. Note that the installation steps have the merged binding information from the new version and the old version. If the new version has bindings for application artifacts such as EJB JNDI names, EJB references or resource references, then those bindings will be part of the merged binding information. If new bindings are not present, then bindings are taken from the installed (old) version. If bindings are not present in the old version and if the default binding generation option is enabled, then the default bindings will be part of the merged binding information. You can select whether to ignore bindings in the old version or ones in the new version.
6. Map the installed application or module to servers or clusters. Use the Map modules to application servers page of the Install New Application pages displayed during updating the application. Or, after updating the application, use the Map modules to application servers page accessed from the Enterprise Applications page.
  - a. Go to the Map modules to application servers page. Click **Applications > Enterprise Applications** in the console navigation tree, click the application name, and then click **Map modules to application servers**.
  - b. Specify the application server where you want to install modules contained in your application and click **OK**.
7. Click **Save** on the administrative console taskbar to save the changes to your configuration. In the single server (base) product, after you click **Save** the old version of the application is uninstalled and the new version is installed into the configuration. The application binaries for the old version are deleted from the destination directory and the new binaries are copied to the directory. In the Network

Deployment product, the old application files are deleted and new files are copied when the configuration on the deployment manager synchronizes with the configuration on the node where the application is installed. If the application is running when you update it, the application stops running before its files are copied to the destination directory of the node and restarts after the copy operation completes. Thus, the application is unavailable on the node during the time the node is synchronizing its configuration with the deployment manager.

8. Examine the values specified for **Reload Enabled** and **Reload Interval** on the settings page for your enterprise application. If reloading of application files is enabled and the reload interval is greater than zero (0), the application's files are reloaded after the application is updated. For Web modules such as servlets and JavaServer page (JSP) files, a Web container reloads a Web module only when the IBM extension reloadingEnabled in the `ibm-web-ext.xmi` file is also set to `true`. You can set reloadingEnabled to `true` when editing your Web module's extended deployment descriptors in an assembly tool.
9. Restart the application so the changes take effect. If the application is updated while it is running, WebSphere Application Server stops the application, updates the application logic and restarts the application.
  - a. Click **Applications > Enterprise Applications** in the console navigation tree to go to the Enterprise Applications page.
  - b. Check the check box for the updated application.
  - c. Click **Start**.
10. Optional: If the application you are updating is deployed on a server that has its application class loader policy set to SINGLE on the Application Server settings page of the administrative console, restart the server.

## Hot deployment and dynamic reloading

You can make various changes to applications and their contents without having to stop the server and start it again. Making these types of changes is known as *hot deployment and dynamic reloading*.

Hot deployment is the process of adding new components (such as WAR files, EJB Jar files, enterprise Java beans, servlets, and JSP files) to a running server without having to stop the application server process and start it again.

Dynamic reloading is the ability to change an existing component without needing to restart the server in order for the change to take effect. Dynamic reloading involves:

- Changes to the implementation of a component of an application, such as changing the implementation of a servlet
- Changes to the settings of the application, such as changing the deployment descriptor for a Web module

If the application you are updating is deployed on a server that has its application class loader policy set to SINGLE on the application server settings page of the administrative console, you might not be able to dynamically reload your application. At minimum, you must restart the server after updating your application.

1. Locate your expanded application files. The application files are in the directory you specified when installing the application or, if you did not specify a custom target directory, are in the default target directory, `install_root/installedApps/cell_name`. Your EAR file, `${APP_INSTALL_ROOT}/cell_name/application_name.ear`, points to the target directory. The `variables.xml` file for the node defines `${APP_INSTALL_ROOT}`.

It is important to locate the expanded application files because, as part of installing applications, a WebSphere application server unjars portions of the EAR file onto the file system of the computer that will run the application. These expanded files are what the server looks at when running your application.

If you cannot locate the expanded application files, look at the `binariesURL` attribute in the `deployment.xml` file for your application. The attribute designates the location the run time uses to find the application files.

For the remainder of this information on hot deployment and dynamic reloading, `application_root` represents the root directory of the expanded application files.

2. Locate application metadata files. The metadata files include the deployment descriptors (`web.xml`, `application.xml`, `ejb-jar.xml`, and the like), the bindings files (`ibm-web-bnd.xmi`, `ibm-app-bnd.xmi`, and the like), and the extensions files (`ibm-web-ext.xmi`, `ibm-app-ext.xmi`, and the like).

Metadata XML files for an application can be loaded from one of two locations. The metadata files can be loaded from the same location as the application binary files (such as `application_root/META-INF`) or they can be loaded from the WebSphere configuration tree,

`${CONFIG_ROOT}/cells/cell_name/applications`

`/application_EAR_name/deployments/application_name/`. The value of the `useMetadataFromBinary` flag specified during application installation controls which location is used. If specified, the metadata files are loaded from the same location as the application binary files. If not specified, the metadata files are loaded from the application deployment folder in the configuration tree.

For the remainder of this information, `metadata_root` represents the location of the metadata files for the specified application or module.

3. CAUTION: If you are running WebSphere Application Server on a group of machines using Network Deployment and you are changing an application on a particular node, disable automatic synchronization.
  - a. Click **System Administration > Node Agents** in the administrative console navigation tree, click on a node agent name, and then click **File Synchronization Service**.
  - b. On the File Synchronization Service page, remove the checkmark from the check box for **Automatic Synchronization** and click **OK**.

When you run WebSphere Application Server on a group of machines using Network Deployment and you change a file on the disk in the expanded application directory for a particular node, you can lose those changes the next time node synchronization occurs. In the Network Deployment environment, the configuration stored by the deployment manager is the master copy and any changes detected between that master copy and the copy on a particular machine trigger the master copy to be downloaded to the node.

4. Optional: Examine the values specified for **Reload Enabled** and **Reload Interval** on the settings page for your enterprise application. If reloading of application files is enabled and the reload interval is greater than zero (0), the application's files are reloaded after the application is updated. For Web modules such as servlets and JavaServer page (JSP) files, a Web container reloads a Web module only when the IBM extension `reloadingEnabled` in the `ibm-web-ext.xmi` file is also set to `true`. You can set `reloadingEnabled` to `true` when editing your Web module's extended deployment descriptors in an assembly tool.
5. Change or add the following components or modules as needed:
  - Application files
  - WAR files
  - EJB Jar files
  - HTTP plug-in configuration files
6. For changes to take effect, you might need to start, stop, or restart an application. "Starting and stopping applications" on page 1421 provides information on using the administrative console to start, stop, or restart an application. "Example: Starting an application using wsadmin" and "Example: Stopping running applications on a server using wsadmin" in the Information Center provide information on using the wsadmin scripting tool.
7. If you disabled automatic synchronization in step 3, return to the File Synchronization Service page, enable **Automatic Synchronization**, and click **OK**.

## Changing or adding application files

You can change or add application files on application servers without having to stop the server and start it again. This file describes--

- Updating an existing application on a running server (providing a new EAR file)
- Adding a new application to a running server
- Removing an existing application from a running server
- Changing or adding files to existing EJB or Web modules
- Changing the `application.xml` file for an application
- Changing the `ibm-app-ext.xmi` file for an application
- Changing the `ibm-app-bnd.xmi` file for an application
- Changing a non-module Jar file contained in the EAR file

### Updating an existing application on a running server (providing a new EAR file)

Reinstall an updated application using the administrative console or the `wsadmin $AdminApp install` command with the `-update` option.

Both reinstallation methods enable you to update an existing application using any of the other steps listed in this file, including changing classes, adding modules, removing modules, changing modules, or changing metadata files. The application reinstallation methods detect the changes in your application and prompt you for additional binding data that might be needed to install the application. The reinstallation process automatically stops and restarts your application on the appropriate servers.

<b>Hot deployment</b>	Yes
<b>Dynamic reloading</b>	Yes

### Adding a new application to a running server

Install an application using the administrative console or the `wsadmin install` command.

<b>Hot deployment</b>	Yes
<b>Dynamic reloading</b>	No

### Removing an existing application from a running server

Stop the application and then uninstall it from the server. Use the administrative console to stop the application and then uninstall it. Or run the `wasadmin stopApplication` command and then the `uninstall` command.

<b>Hot deployment</b>	Yes
<b>Dynamic reloading</b>	No

### Changing or adding files to existing EJB or Web modules

1. Update the application files in the `application_root` location.
2. Restart the application. Use the administrative console to restart the application. Or run the `wasadmin stopApplication` and `startApplication` commands.

<b>Hot deployment</b>	Yes
<b>Dynamic reloading</b>	No

### Changing the `application.xml` file for an application

Restart the application. Automatic reloading will not detect the change. Use the administrative console to restart the application. Or run the wasadmin stopApplication and startApplication commands.

<b>Hot deployment</b>	Not applicable
<b>Dynamic reloading</b>	Yes

### Changing the ibm-app-ext.xmi file for an application

Restart the application. Automatic reloading will not detect the change. Use the administrative console to restart the application. Or run the wasadmin stopApplication and startApplication commands.

<b>Hot deployment</b>	Not applicable
<b>Dynamic reloading</b>	Yes

### Changing the ibm-app-bnd.xmi file for an application

Restart the application. Automatic reloading will not detect the change. Use the administrative console to restart the application. Or run the wasadmin stopApplication and startApplication commands.

<b>Hot deployment</b>	Not applicable
<b>Dynamic reloading</b>	Yes

### Changing a non-module Jar file contained in the EAR file

1. Update the non-module Jar file in the *application\_root* location.
2. If automatic reloading is not enabled, restart the application. Use the administrative console to restart the application. Or run the wasadmin stopApplication and startApplication commands.  
If automatic reloading is enabled, you do not need to take further action. Automatic reloading will detect the change.

<b>Hot deployment</b>	Yes
<b>Dynamic reloading</b>	Yes

### Changing or adding WAR files

You can change WAR files on application servers without having to stop the server and start it again. This file describes--

- Changing an existing JSP file
- Adding a new JSP file to an existing application
- Changing an existing servlet class (editing and recompiling)
- Changing a dependent class of an existing servlet class
- Adding a new servlet using the Invoker (Serve Servlets by class name) facility or adding a dependent class to an existing application
- Adding a new servlet, including a new definition of the servlet in the web.xml deployment descriptor for the application
- Changing the web.xml file of a WAR file
- Changing the ibm-web-ext.xmi file of a WAR file
- Changing the ibm-web-bnd.xmi file of a WAR file

### Changing an existing JSP file

Place the changed JSP file directly in the *application\_root/module\_name* directory or the appropriate subdirectory. The change will be automatically detected and the JSP will be recompiled and reloaded.

<b>Hot deployment</b>	Not applicable
-----------------------	----------------



**Dynamic reloading** Yes

### Adding a new JSP file to an existing application

Place the new JSP file directly in the *application\_root/module\_name* directory or the appropriate subdirectory. The new file will be automatically detected and compiled on the first request to the page.

**Hot deployment** Yes

**Dynamic reloading** Yes

### Changing an existing servlet class (editing and recompiling)

1. Place the new version of the servlet `.class` file directly in the *application\_root/module\_name/WEB-INF/classes* directory. If the `.class` file is part of a Jar file, you can place the new version of the Jar file directly in *application\_root/module\_name/WEB-INF/lib*. In either case, the change will be detected, the Web application will be shut down and reinitialized, picking up the new class.
2. If automatic reloading is not enabled, restart the application. Use the administrative console to restart the application. Or run the `wasadmin stopApplication` and `startApplication` commands.

If automatic reloading is enabled, you do not need to take further action. Automatic reloading will detect the change.

**Hot deployment** Not applicable

**Dynamic reloading** Yes

### Changing a dependent class of an existing servlet class

1. Place the new version of the dependent `.class` file directly in the *application\_root/module\_name/WEB-INF/classes* directory. If the `.class` file is part of a Jar file, you can place the new version of the Jar file directly in *application\_root/module\_name/WEB-INF/lib*. In either case, the change will be detected, the Web application will be shut down and reinitialized, picking up the new class.
2. If automatic reloading is not enabled, restart the application. Use the administrative console to restart the application. Or run the `wasadmin stopApplication` and `startApplication` commands.

If automatic reloading is enabled, you do not need to take further action. Automatic reloading will detect the change.

**Hot deployment** Not applicable

**Dynamic reloading** Yes

### Adding a new servlet using the Invoker (Serve Servlets by class name) facility or adding a dependent class to an existing application

1. Place the new `.class` file directly in the *application\_root/module\_name/WEB-INF/classes* directory. If the `.class` file is part of a Jar file, you can place the new version of the Jar file directly in *application\_root/module\_name/WEB-INF/lib*. In either case, the change will be detected, the Web application will be shut down and reinitialized, picking up the new class.

This case is treated the same as changing an existing class. The difference is that adding the servlet or class does not immediately cause the Web application to reload because the class has never been loaded before. The class simply becomes available for execution.

2. If automatic reloading is not enabled, restart the application. Use the administrative console to restart the application. Or run the `wasadmin stopApplication` and `startApplication` commands.

If automatic reloading is enabled, you do not need to take further action. Automatic reloading will detect the change.

**Hot deployment** Yes

**Dynamic reloading** Not applicable

### **Adding a new servlet, including a new definition of the servlet in the web.xml deployment descriptor for the application**

1. Place the new .class file directly in the *application\_root/module\_name/WEB-INF/classes* directory. If the “.class” file is part of a Jar file, you can place the new version of the Jar file directly in *application\_root/module\_name/WEB-INF/lib*.

You can edit the web.xml file in place or copy it into the *application\_root/module\_name/WEB-INF/classes* directory. The new .class file will not trigger a reloading of the application.

2. Restart the application. Use the administrative console to restart the application. Or run the wasadmin stopApplication and startApplication commands. After the application restarts, the new servlet is available for service.

<b>Hot deployment</b>	Yes
<b>Dynamic reloading</b>	Not applicable

### **Changing the web.xml file of a WAR file**

1. Edit the web.xml file in place or copy it into the *metadata\_root/module\_name/WEB-INF* directory.
2. Restart the application. Use the administrative console to restart the application. Or run the wasadmin stopApplication and startApplication commands.

<b>Hot deployment</b>	Yes
<b>Dynamic reloading</b>	Yes

### **Changing the ibm-web-ext.xmi file of a WAR file**

Edit the extension settings as needed. You can change all of the extension settings. The only warning is if you set the reloadInterval property to zero (0) or the reloadEnabled property to false, the application will no longer automatically detect changes to class files. Both of these changes disable the automatic reloading function. The only way to re-enable automatic reloading is to change the appropriate property and restart the application. See other task descriptions in this file for information on restarting an application.

<b>Hot deployment</b>	Not applicable
<b>Dynamic reloading</b>	Yes

### **Changing the ibm-web-bnd.xmi file of a WAR file**

1. Edit the bindings as needed. You can change all of the values but ensure that the entities you are binding to are present in the configuration of the server.
2. Restart the application. Use the administrative console to restart the application. Or run the wasadmin stopApplication and startApplication commands.

<b>Hot deployment</b>	Not applicable
<b>Dynamic reloading</b>	Yes

### **Changing or adding EJB Jar files**

You can change EJB Jar files on application servers without having to stop the server and start it again. This file describes--

- Changing the ejb-jar.xml file of an EJB Jar file
- Changing the ibm-ejb-jar-ext.xmi or ibm-ejb-jar-bnd.xmi file of an EJB Jar file
- Changing the Table.ddl file for an EJB Jar file
- Changing the Map.mapxmi or Schema.dbxmi file for an EJB Jar file
- Updating the implementation class for an EJB file or a dependent class of the implementation class for an EJB file
- Updating the Home/Remote interface class for an EJB file
- Adding a new EJB file to an existing EJB Jar file

### Changing the `ejb-jar.xml` file of an EJB Jar file

Restart the application. Automatic reloading will not detect the change. Use the administrative console to restart the application. Or run the `wasadmin stopApplication` and `startApplication` commands.

<b>Hot deployment</b>	Not applicable
<b>Dynamic reloading</b>	Yes

### Change the `ibm-ejb-jar-ext.xml` or `ibm-ejb-jar-bnd.xml` file of an EJB Jar file

Restart the application. Automatic reloading will not detect the change. Use the administrative console to restart the application. Or run the `wasadmin stopApplication` and `startApplication` commands.

<b>Hot deployment</b>	Not applicable
<b>Dynamic reloading</b>	Yes

### Changing the `Table.ddl` file for an EJB Jar file

Rerun the DDL file on the user database server. Changing the `Table.ddl` file has no effect on the application server and is a change to the database table schema for the EJB files.

<b>Hot deployment</b>	Not applicable
<b>Dynamic reloading</b>	Not applicable

### Changing the `Map.mapxml` or `Schema.dbxml` file for an EJB Jar file

1. Change the `Map.mapxml` or `Schema.dbxml` file for an EJB Jar file.
2. Regenerate the deployed code artifacts for the EJB file.
3. Apply the new EJB Jar file to the server.
4. Restart the application. Use the administrative console to restart the application. Or run the `wasadmin stopApplication` and `startApplication` commands.

<b>Hot deployment</b>	Not applicable
<b>Dynamic reloading</b>	Yes

### Updating the implementation class for an EJB file or a dependent class of the implementation class for an EJB file

1. Update the class file in the `application_root/module_name.jar` file.
2. If automatic reloading is enabled, you do not need to take further action. Automatic reloading will detect the change.

If automatic reloading is not enabled, restart the application of which the EJB file is a member. If the updated module is used by other modules in other applications, restart those applications as well. Use the administrative console to restart the application. Or run the `wasadmin stopApplication` and `startApplication` commands.

<b>Hot deployment</b>	Not applicable
<b>Dynamic reloading</b>	Yes

### Updating the Home/Remote interface class for an EJB file

1. Update the interface class of the EJB file.
2. Regenerate the deployed code artifacts for the EJB file.
3. Apply the new EJB Jar file to the server.
4. If automatic reloading is enabled, you do not need to take further action. Automatic reloading will detect the change.

If automatic reloading is not enabled, restart the application of which the EJB file is a member. Use the administrative console to restart the application. Or run the wasadmin stopApplication and startApplication commands.

<b>Hot deployment</b>	Not applicable
<b>Dynamic reloading</b>	Yes

### Adding a new EJB file to an existing EJB Jar file

1. Apply the new or updated Jar file to the *application\_root* location.
2. If automatic reloading is enabled, you do not need to take further action. Automatic reloading will detect the change.

If automatic reloading is not enabled, restart the application. Use the administrative console to restart the application. Or run the wasadmin stopApplication and startApplication commands.

<b>Hot deployment</b>	Yes
<b>Dynamic reloading</b>	Yes

## Changing the HTTP plug-in configuration

You can change the HTTP plug-in configuration without having to stop the server and start it again. This file describes--

- Changing the application.xml file to change the context root of a WAR file
- Changing the web.xml file to add, remove, or modify a servlet mapping
- Changing the server.xml file to add, remove, or modify an HTTP transport or changing the virtualhost.xml file to add or remove a virtual host or to add, remove, or modify a virtual host alias

### Changing the application.xml file to change the context root of a WAR file

1. Change the application.xml file.
2. Regenerate the plug-in configuration file using the administrative console or by running the GenPluginCfg.bat/sh script.

<b>Hot deployment</b>	Yes
<b>Dynamic reloading</b>	No

### Changing the web.xml file to add, remove, or modify a servlet mapping

1. Change the web.xml file.
2. Regenerate the plug-in configuration file using the administrative console or by running the GenPluginCfg.bat/sh script.

If the Web application has file serving enabled or has a servlet mapping of /, you do not have to regenerate the plug-in configuration. In all other cases the regeneration is required.

<b>Hot deployment</b>	Yes
<b>Dynamic reloading</b>	Yes

### Changing the server.xml file to add, remove, or modify an HTTP transport or changing the virtualhost.xml file to add or remove a virtual host or to add, remove, or modify a virtual host alias

1. Change the server.xml file to add, remove, or modify an HTTP transport or change the virtualhost.xml file to add or remove a virtual host or to add, remove, or modify a virtual host alias.
2. Regenerate the plug-in configuration file using the administrative console, by running the GenPluginCfg.bat/sh script, or by running a wsadmin command.

<b>Hot deployment</b>	Yes
<b>Dynamic reloading</b>	Yes

---

## Uninstalling applications

After an application no longer is needed, you can uninstall it. Uninstalling an application deletes the application from the WebSphere Application Server configuration repository and it deletes the application binaries from the file system of all nodes where the application modules are installed.

1. Click **Applications > Enterprise Applications** in the administrative console navigation tree to access the Enterprise Applications page.
2. Stop the application. Select the application you want uninstalled and click **Stop**.
3. Back up the application. Select the application you want uninstalled and click **Export** to export the application to an EAR file and preserve the binding information.
4. Select the application you want uninstalled and click **Uninstall**.
5. Click **Save** on the console taskbar to save changes made to the administrative configuration.

In the single-server (base) product, application binaries are deleted after you click **Save**. In the Network Deployment product, application binaries are deleted when configuration changes on the deployment manager synchronize with configurations for individual nodes.

---

## Deploying and managing applications: Resources for learning

Use the following links to find relevant supplemental information about deploying and managing applications using the administrative console. The information resides on IBM and non-IBM Internet sites, whose sponsors control the technical accuracy of the information.

These links are provided for convenience. Often, the information is not specific to the IBM WebSphere Application Server product, but is useful all or in part for understanding the product. When possible, links are provided to technical papers and Redbooks that supplement the broad coverage of the release documentation with in-depth examinations of particular product areas.

View links to additional information about:

- Programming model and decisions
- Programming instructions and examples
- Programming instructions and examples

### Programming model and decisions

- The J2EE™ Tutorial: The Duke's Bank Application
- Best Practices in WebSphere Application: Separating the developers from the administrators
- Designing Enterprise Applications with the Java™ 2 Platform, Enterprise Edition, Second Edition
- Designing Enterprise Applications, Second Edition
- Building Java™ Enterprise Applications Volume I: Architecture

### Programming instructions and examples

- WebSphere Application Server education
- Developing and Testing a Complete 'Hello World' J2EE Application with IBM WebSphere Studio Application Developer for Linux
- Writing Enterprise Applications with Java™ 2 Platform, Enterprise Edition

### Administration

- Listing of all IBM WebSphere Application Server Redbooks



---

## Notices

References in this publication to IBM products, programs, or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM product, program, or service is not intended to state or imply that only IBM's product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any of IBM's intellectual property rights may be used instead of the IBM product, program, or service. Evaluation and verification of operation in conjunction with other products, except those expressly designated by IBM, is the user's responsibility.

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing  
IBM Corporation  
500 Columbus Avenue  
Thornwood, New York 10594 USA





---

## Trademarks and service marks

The following terms are trademarks of IBM Corporation in the United States, other countries, or both:

- AIX
- AS/400
- CICS
- Cloudscape
- DB2
- DFSMS
- Domino
- Everyplace
- iSeries
- IBM
- IMS
- Informix
- iSeries
- Language Environment
- Lotus
- MQSeries
- MVS
- OS/390
- RACF
- Redbooks
- RMF
- SecureWay
- SupportPac
- Tivoli
- ViaVoice
- VisualAge
- VTAM
- WebSphere
- z/OS
- zSeries

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Other company, product, and service names may be trademarks or service marks of others.