

IBM WebSphere Application Server for z/OS, Version 8.0

*Administering applications and their
environment*



Note

Before using this information, be sure to read the general information under “Notices” on page 607.

Compilation date: June 9, 2011

© Copyright IBM Corporation 2011.

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

Contents

How to send your comments.	ix
Changes to serve you more quickly	xi
Chapter 1. Overview and new features for administering applications and their environments	1
Chapter 2. How do I administer applications and their environments?	3
Chapter 3. Using the administrative clients	5
Using the administrative console	5
Administrative console	5
Installing and uninstalling the administrative console	19
Starting and logging off the administrative console	20
Specifying console preferences	23
Accessing help and product information from the administrative console	29
Changing the console session expiration	54
Changing the class loader order of the console module deployed in Integrated Solutions Console	55
Getting started with wsadmin scripting	56
What is new for scripted administration (wsadmin)	58
Overview and new features for scripting the application serving environment.	59
Using administrative programs (JMX)	59
Java Management Extensions (JMX) for WebSphere Application Server	61
Creating a custom Java administrative client program using WebSphere Application Server administrative Java APIs	63
Creating a Java Management Extensions client program using the Java Management Extensions Remote application programming interface	75
Extending the WebSphere Application Server administrative system with custom MBeans	79
Administrative programs for multiple Java Platform, Enterprise Edition application servers	104
Deploying and managing a custom Java administrative client program with multiple Java Platform, Enterprise Edition application servers	106
Java Management Extensions V1.0 to Java Management Extensions V1.2 migration	107
Java Management Extensions (JMX) interoperability	108
Managed object metadata	108
Using command-line tools	109
startServer command	110
stopServer command.	113
startManager command.	115
stopManager command.	116
startNode command	117
stopNode command	118
serverStatus command	120
addNode command	121
addNode command best practices	127
removeNode command	128
cleanupNode command.	130
syncNode command	131
renameNode command.	132
registerNode command.	133
deregisterNode command	136
backupConfig command	137
restoreConfig command	139
versionInfo command	140
Location of the command file.	141

Syntax for the versionInfo command	141
Parameters	141
Report description	141
Sample versionInfo report	142
genVersionReport command	143
Location of the command file	143
Syntax for the genVersionReport command	143
Report description	143
Sample genVersionReport report	145
historyInfo command	145
Location of the command file	145
Syntax for the historyInfo command	145
Parameters	146
Report description	146
Sample historyInfo report	146
genHistoryReport command	147
Location of the command file	148
Syntax for the genHistoryReport command	148
Report description	148
Sample historyInfo report	148
managesdk command	149
GenPluginCfg command	153
EARExpander command	155
revokeCertificate command	156
requestCertificate command	158
createCertRequest command	160
queryCertificate command	161
restoreJobManager command	163
Example: Security and the command line tools	163
Using MVS console commands	164
START command	164
STOP command	165
Modify command	165
Display command with examples	178
Example: Displaying units of work (transactions) for the Information Management System	186
Example: Displaying the setting for trace records	186
Chapter 4. Using Ant to automate tasks	187
Chapter 5. Starting and stopping quick reference	189
Chapter 6. Backing up and recovering the application serving environment.	191
Chapter 7. Class loading	193
Class loaders	193
Configuring class loaders of a server	198
Class loader collection	199
Class loader ID	199
Class loader order	199
Class loader settings	199
Configuring application class loaders	200
Configuring web module class loaders	202
Class loading: Resources for learning	203
Chapter 8. Deploying and administering enterprise applications	205
Enterprise (Java EE) applications	206

System applications	206
Common deployment framework	206
Installing enterprise application files	207
Installable enterprise module versions	208
Ways to install enterprise applications or modules	210
Installing enterprise application files with the console	213
Example: Installing an EAR file using the default bindings	220
Example: Installing a web services sample with the console	221
Preparing for application installation settings	222
Preparing for application installation binding settings	223
Select installation options settings	228
Manage modules settings	238
Client module settings	241
Client module property settings	241
Provide options to compile JavaServer Pages settings	241
EJB JNDI names for beans	243
Bind EJB business settings	244
Map default data sources for modules containing 1.x entity beans	244
EJB references	245
Resource references	246
Virtual hosts settings	249
Security role to user or group mapping	250
JASPI authentication enablement for applications	251
User RunAs collection	252
Ensure all unprotected 1.x methods have the correct level of protection	253
Bind listeners for message-driven beans settings	253
Map data sources for all 2.x CMP beans	255
Map data sources for all 2.x CMP beans settings	257
Ensure all unprotected 2.x methods have the correct level of protection	259
Provide options to perform the EJB Deploy settings	260
Shared library reference and mapping settings	262
Shared library relationship and mapping settings	263
JSP and JSF option settings	264
Context root for web modules settings	266
Initial parameters for servlets settings	266
Environment entries for client modules settings	267
Environment entries for EJB modules settings	268
Environment entries for web modules settings	268
Environment entries for application settings	269
Resource environment references	269
Message destination reference settings	270
Select current backend ID settings	271
Provide JNDI names for JCA objects settings	272
Correct use of the system identity	272
Requirements for setting data access isolation levels	273
Metadata for module settings	275
Provide options to perform the web services deployment settings	277
Display module build ID settings	278
Task overview: Assembling applications using remote request dispatcher	278
Installing enterprise modules with JSR-88	280
Customizing modules using DConfigBeans	282
Configuring enterprise application files	283
Application bindings	284
Enterprise application collection	290
Configuring application startup	293
Configuring binary location and use	295

Configuring the use of class loaders by an application	301
Manage modules settings	305
Mapping modules to servers	307
Mapping virtual hosts for web modules	309
Mapping properties for a custom login or trusted connection configuration	312
Viewing deployment descriptors	312
Metadata for module settings	315
Starting or stopping enterprise applications	316
Disabling automatic starting of applications	318
Target specific application status	319
Updating enterprise application files	320
Ways to update enterprise application files	322
Updating enterprise applications with the console	324
Preparing for application update settings	326
Hot deployment and dynamic reloading	330
Resolving application configuration conflicts	340
Exporting enterprise applications	342
Exporting enterprise application files	344
Exporting DDL files	344
Uninstalling enterprise applications using the console	345
Removing enterprise files	346
Uninstalling enterprise application files by dragging them from a monitored directory	347
Deploying and administering applications: Resources for learning	349
Chapter 9. Managing applications through programming	351
Accessing the application management function	352
Installing an application through programming	353
Application management	356
Starting an application through programming	358
Sharing sessions for application management	359
Manipulating additional attributes for a deployed application	360
Editing applications	362
Updating an application through programming	363
Adding to, updating, or deleting part of an application through programming	366
Preparing a module and adding it to an existing application through programming	367
Preparing and updating a module through programming	370
Adding a file through programming	372
Updating a file through programming	374
Uninstalling an application through programming	376
Deleting a module through programming	378
Deleting a file through programming	380
Chapter 10. Extending application management operations through programming	383
Chapter 11. Deploying and administering business-level applications	387
Business-level applications	388
Assets	390
Composition units	391
Importing assets	392
Upload asset settings	394
Asset settings	394
Managing assets	398
Asset collection	398
Updating assets	399
Deleting assets	403
Exporting assets	404

Creating business-level applications	404
Creating business-level applications with the console	405
Business-level application settings	415
Composition unit settings	417
Example: Creating a business-level application	420
SCA application package deployment	421
Creating SCA business-level applications	423
Creating SCA business-level applications with the console	425
Composition unit settings	433
JMS binding settings for SCA composites	436
Provide HTTP endpoint URL information settings for SCA composites.	441
SCA composite component settings	441
SCA component reference settings	442
SCA component service settings	443
Service provider policy sets and bindings collection for SCA composites	443
References policy sets and bindings collection for SCA composites	446
SCA service provider settings	448
SCA service client settings	452
Example: Creating an SCA business-level application with the console	455
Starting business-level applications	457
Stopping business-level applications	458
Updating business-level applications	458
Updating SCA composite artifacts	460
Viewing SCA composite definitions	461
Viewing SCA domain information	462
Viewing and editing JMS bindings on references and services of SCA composites	463
Exporting WSDL and XSD documents	464
Deleting business-level applications	465
Chapter 12. Administering business-level applications using programming	469
Creating an empty business-level application using programming	471
Importing an asset using programming	474
Listing assets using programming	479
Viewing an asset using programming.	483
Editing an asset using programming	487
Deleting an asset using programming	491
Exporting an asset using programming	495
Starting a business-level application using programming.	498
Stopping a business-level application using programming	502
Checking the status of a business-level application using programming	505
Listing business-level applications using programming	509
Listing composition units using programming	512
Listing control operations using programming.	516
Viewing a business-level application using programming.	520
Viewing a composition unit using programming	524
Adding a composition unit using programming	528
Updating an asset using programming	534
Editing a business-level application using programming	538
Editing a composition unit using programming	542
Deleting a business-level application using programming	548
Deleting a composition unit using programming	551
Chapter 13. Troubleshooting deployment	557
Application deployment problems	557
Application deployment troubleshooting tips	563
Application startup errors	563

Application startup problems	568
Reducing annotation searches during application deployment	570
A client program does not work	571
Web resource is not displayed	572
Application uninstallation problems.	574
Chapter 14. Troubleshooting administration	577
Administration and administrative console troubleshooting	577
Administrative console does not start even though installation completes	581
Administrative console - browser connection problems	582
When a single user that uses multiple instances of the Mozilla browser logs into the administrative console, the first user ID that logs into the administrative console is the current user.	582
A user on Mozilla browser Version 1.4 selects a check box on a collection table, presses Enter, and receives an error.	583
A user on Mozilla browser Version 1.4 enters an invalid ID or password, presses Enter, and receives an error message.	583
Web server plug-in troubleshooting tips	583
Multiserver environment errors	585
Workload management component troubleshooting tips	588
Workload is not getting distributed	589
Administrative problems with the wsadmin scripting tool	591
Tracing and logging facilities - troubleshooting tips	597
Cannot restart the Deployment Manager monitoring policy	597
Timeout condition resolutions.	597
Timer overview	597
Timeout conditions: analyzing diagnostic data	599
Timeout conditions - possible causes and fixes	601
Timeout values: guidelines for altering timeout values.	602
Appendix. Directory conventions	605
Notices	607
Trademarks and service marks	609
Index	611

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 919-254-5250.

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.

Changes to serve you more quickly

Print sections directly from the information center navigation

PDF books are provided as a convenience format for easy printing, reading, and offline use. The information center is the official delivery format for IBM WebSphere Application Server documentation. If you use the PDF books primarily for convenient printing, it is now easier to print various parts of the information center as needed, quickly and directly from the information center navigation tree.

To print a section of the information center navigation:

1. Hover your cursor over an entry in the information center navigation until the **Open Quick Menu** icon is displayed beside the entry.
2. Right-click the icon to display a menu for printing or searching your selected section of the navigation tree.
3. If you select **Print this topic and subtopics** from the menu, the selected section is launched in a separate browser window as one HTML file. The HTML file includes each of the topics in the section, with a table of contents at the top.
4. Print the HTML file.

For performance reasons, the number of topics you can print at one time is limited. You are notified if your selection contains too many topics. If the current limit is too restrictive, use the feedback link to suggest a preferable limit. The feedback link is available at the end of most information center pages.

Under construction!

The Information Development Team for IBM WebSphere Application Server is changing its PDF book delivery strategy to respond better to user needs. The intention is to deliver the content to you in PDF format more frequently. During a temporary transition phase, you might experience broken links. During the transition phase, expect the following link behavior:

- Links to Web addresses beginning with `http://` work
- Links that refer to specific page numbers within the same PDF book work
- The remaining links will *not* work. You receive an error message when you click them

Thanks for your patience, in the short term, to facilitate the transition to more frequent PDF book updates.

Chapter 1. Overview and new features for administering applications and their environments

Use the links provided in this topic to learn about the administrative features.

What is new for administrators

This topic provides an overview of new and changed features of system administration.

Introduction: System administration

This topic describes the administration of the product and the applications that run on it.

See also Introduction: Environment and Introduction: Variables.

Chapter 2. How do I administer applications and their environments?

Follow these shortcuts to get started quickly with popular tasks.

When you visit a task in the information center, look for the **IBM Suggests** feature at the bottom of the page. Use it to find available tutorials, demonstrations, presentations, developerWorks® articles, IBM® Redbooks®, support documents, and more.

Administer nodes

Administer node agents

Administer cells

Register a node with the administrative agent

Administer nodes using the administrative agent

Register a node with the job manager

Administer nodes from the console using the job manager

Administer nodes from the wsadmin tool using the job manager

Administer configurations

Configure remote file services

Administer application servers with the console

Configure application servers with scripting

Manage application servers with scripting

Administer generic servers

Use server templates

Administer custom services

Use the console to balance workloads by clustering application servers

Use scripting to balance workloads by clustering application server

Establish high availability (HA)

Administer the UDDI registry

Use the console to administer communication with web servers (plug-ins)

Use scripting to administer communication with web servers (plug-ins)

Administer HTTP sessions with the console

Administer HTTP sessions with scripting

Provide access to naming and directory resources (JNDI) - Name server

Provide access to naming and directory resources (JNDI) - Bindings

Provide access to relational databases (JDBC resources) with the console

Provide access to relational databases (JDBC resources) with scripting

Choose a messaging provider.

Provide access to messaging resources (default messaging provider) with scripting

Install applications with the console

Install applications with scripting

Start and stop applications with the console

Start and stop applications with scripting

Update applications with the console

Update applications with scripting

Deploy and administer web services applications

Administer business-level applications using the administrative console

Administer business-level applications using programming

Set up business-level applications using scripting

Manage environment configurations with properties files using wsadmin scripting

Choose an administrative client

Use the administrative console

Using scripting (wsadmin)

Troubleshoot deployment

Troubleshoot administration

Chapter 3. Using the administrative clients

Using the administrative console

You can install, start, and access the administrative console. You can also specify console preferences and access help.

About this task

The administrative console is a Web-based tool that you use to manage the product. The administrative console supports a full range of product administrative activities.

Unfederated application servers, administrative agents, deployment managers, and job managers can have their own administrative consoles. The steps in this task apply to these consoles. For federated application servers, the administrative console belongs to the deployment manager.





Procedure

1. Optionally install the administrative console through the wsadmin command.
2. Start the server for the appropriate administrative console.
Issue the startServer command or the MVS™ console START command for an unfederated application server, an administrative agent, or a job manager. Issue the StartManager command or the MVS console START command for a deployment manager. The administrative console application starts automatically when you start the server to which the administrative console belongs.
3. Access the administrative console.
4. Specify console preferences.
5. Access help.

Administrative console

Administrative console buttons

This page describes the button choices that are available on various pages of the administrative console, depending on which product features you enable.

-  **Select all items.** Selects each resource that is listed on the administrative console panel, in preparation for performing an action against the selected resources.
-  **Deselect all items.** Removes all the listed resources from each selection so that no action is performed against any of the resources.
-  **Show filter function.** Produces a dialog box for specifying the resources to view in the table on this administrative console page.
-  **Hide filter function.** Hides the dialog box for specifying the resources to view in the table on this administrative console page.

When you produce the dialog box, select the column to filter and enter the filter criteria.

Column to filter

Select the column to filter from the drop-down list. When you apply the filter, only those items in the selected column that meet the filter criteria are displayed.

For example, select **Names** to enter criteria by which to filter application server names.

Filter criteria

Enter a string that must be found in the name of a collection entry to qualify the entry to display

in the collection table. The string can contain percent sign (%), asterisk (*), or question mark (?) symbols as wildcard characters. For example, enter *App* to find any application server whose name contains the string App.

Prefix each of the following characters () ^ * % { } \ + \$ with a backslash (\) so that the regular expression engine performing the search correctly matches the search criteria. For example, to search for all Java DataBase Connectivity (JDBC) providers containing (XA) in the provider name, specify the following string:

*\ (XA\)



- **Clear filter value.** Clears your filter changes and restores the most recently saved values.
- **Abort.** Stops a transaction that is not yet in the prepared state. All operations that the transaction completed are undone.
- **Activate.** Activates a group member.
- **Add.** Adds the selected or typed item to a list, or produces a dialog for adding an item to a list.
- **Add Node.** Displays the Add Node page, in which you specify the host name and SOAP connector port for a node that you want added to a cell.
- **Apply.** Saves your changes to a page without exiting the page.
- **Assign Binding.** Assigns a binding to a policy set from a list of available bindings.
- **Attach.** Attaches a policy set to your application, service provider, endpoint or operation.
- **Back.** Displays the previous page or item in a sequence. The administrative console does not support using the Back and Forward options of a browser, which can cause intermittent problems. Use Back or Cancel on the administrative console panels instead.
- **Balance.** Balances active members in high availability groups across servers that host the high availability groups. The administrator must first determine which groups have active members and select those groups before selecting Balance.
- **Browse.** Opens a dialog that enables you to look for a file on your system.
- **Calculate groups.** Calculates the number of high availability groups that are returned based on the match set.
- **Cancel.** Exits the current page or dialog, discarding unsaved changes. The administrative console does not support using the Back and Forward options of a browser, which can cause intermittent problems. Use Cancel on the administrative console panels instead.
- **Change.** In the context of security, you can search the user registry for a user ID for an application to run under. In the context of container properties, you can change the data source that the container is using.
- **Clear.** Clears your changes and restores the most recently saved values.
- **Clear selections.** Clears any selected cells in the tables on this tabbed page.
- **Close.** Exits the dialog.
- **Collapse all.** Collapses all the expanded items.
- **Commit.** Releases all locks that are held by a prepared transaction and forces the transaction to commit.
- **Copy.** Creates copies of the selected application servers.
- **Create.** Saves your changes to all the tabbed pages in a dialog and exits the dialog.
- **Create tables.** Develops scheduler database tables.
- **Deactivate.** Deactivates a group member. The group member must be in the active state to be deactivated. The deactivate option causes the group member to move to the idle state. The group policy overrides which members are activated and deactivated for a group. The policy is enforced for every member state change. If the deactivate option conflicts with the group policy, the policy resets who is the active member of the group.
- **Delete.** Removes the selected instance.
- **Detach.** Detaches a policy set from an application, service provider, endpoint or operation.
- **Details.** Shows the details about a transaction.
- **Disable.**
Group or group members:
Disables a group or group member. When you disable a group or group member, the active

group or group member is first deactivated. If the deactivate option is successful, the group or group member moves to the disable state. A disabled group or group member cannot be activated.

Java Management Extensions (JMX) connectors:

Disables a connector.

- **Disable Auto Start.** Requires you to start the application manually.
- **Discard.** Discards your local changes instead of saving them to the master configuration.
- **Done.** Saves your changes to all the tabbed pages in a dialog and exits the dialog.
- **Down.** Moves through a list.
- **Drop tables.** Removes scheduler database tables.
- **Dump.** Activates a dump of a traced application server.
- **Edit.** Lets you edit the selected item in a list, or produce a dialog box for editing the item.
- **Enable.**

Group or group members:

Enables a group or a group member.

Java Management Extensions (JMX) connectors:

Enables a connector.

- **Enable Auto Start.** Starts an application automatically when the server on which the application resides starts.
- **Expand all.** Expands all the collapsed items.
- **Export.** Accesses a page for exporting enterprise archive (EAR) files for an enterprise application.
- **Export DDL.** Accesses a page for exporting data definition language (DDL) files for an enterprise application.
- **Export Keys.** Exports Lightweight Third-Party Authentication (LTPA) keys to other domains.
- **Export route table.** Exports the route table information for a selected cluster to a binary file in the configuration.
- **Filter.** Produces a dialog box for specifying the resources to view in the tables on this tabbed page.
- **Finish.** Forces a transaction to finish, regardless of whether its outcome has been reported to all participating applications.
- **First.** Displays the first record in a series of records.
- **Force delete.** Forces the removal of a node that is not removed properly from the cell in the master repository. The **Remove node** action is preferred over the **Force delete** action to delete a node from the configuration. If you click **Force delete**, but the node still exists in the configuration, uninstall the node or run the **removeNode** command by using the **-force** parameter on that node. Force delete action is equivalent to running the **cleanupNode** command at the deployment manager.
- **Full resynchronize.** Synchronizes the user's configuration immediately. Click full resynchronize on the Nodes page if automatic configuration synchronization is disabled, or if the synchronization interval is set to a long time, and a configuration change is made to the cell repository that needs to be replicated to that node. Clicking this option clears all synchronization optimization settings and performs configuration synchronization again, so no mismatches occur between node and cell configuration after this operation is performed. This operation can take awhile to perform.
- **Generate keys.** Generates new LTPA keys. When security is turned on for the first time with LTPA as the authentication mechanism, LTPA keys are automatically generated with the password entered in the panel. To generate new keys, use this option after the server is up with security turned on. Clicking this option generates the keys and propagates them to all active servers (cell, node, and application servers). The new keys can be used to encrypt and decrypt the LTPA tokens. Click **Save** on the console taskbar to save the new keys and the password in the repository.
- **Immediate stop.** Stops the server, but bypasses the normal server quiesce process that supports in-flight requests to complete before shutting down the entire server process. This shutdown mode is faster than the normal server stop processing, but some application clients can receive exceptions.
- **Import keys.** Imports new LTPA keys from other domains. To support single sign-on (SSO) in WebSphere® Application Server across multiple WebSphere domains (cells), share LTPA keys and a password among the domains. After exporting the keys from one of the cells into a file, click this option to import the keys into all the active servers (cell, node, and application servers). The new keys can be used to encrypt and decrypt the LTPA token. Click **Save** on the console taskbar to save the new keys and the password in the repository.

- **Install.** Displays the Preparing for application installation page, which you use to deploy an application, an enterprise bean, or a web component onto an application server.
- **Install RAR.** Opens a dialog that is used to install a Java 2 Platform, Enterprise Edition Connector Architecture (JCA) connector and to create a resource adapter.
- **Manage state.** Displays a list of MBeans that corresponds to your previous selection of data source or connection factory configurations. You can apply JCA lifecycle management operations to these MBeans to control the runtime status of the corresponding resources.
- **Manage transactions.** Displays a list of active transactions running on a server. You can forcibly finish any transaction that has stopped processing because a transactional resource is not available.
- **Modify.** Opens a dialog that is used to change a specification.
- **Move.** Moves the selected application servers to a different location in the administrative cell. When prompted, specify the target location.
- **Move down.** Moves downward through a list.
- **Move up.** Moves upward through a list.
- **New.** Displays a page that you use to define a new instance. For example, clicking **New** on the Application Servers page displays a page on which you can configure a new application server.
- **Next.** Displays the next page, frame, or item in a sequence.
- **OK.** Saves your changes to the local configuration and exits the page.
- **Pause.** In the context of JCA lifecycle management, stops all outbound communication that is conducted through a resource on a specified server to a backend.
- **Ping.** Attempts to contact selected application servers.
- **Previous.** Displays the previous page, frame, or item in a sequence.
- **Quit.** Exits a dialog box and discards any unsaved changes.
- **Reference shared libraries.** Opens the collection of shared library references available for use by your application or module. If no references are available, a message is displayed stating that there are no references.
- **Refresh.** Refreshes the view of data for instances that are currently listed on this tabbed page.
- **Remove.** Deletes the selected item.
- **Remove file.** Removes the specified file from the selected application or module.
- **Remove node.** Deletes the selected node.
- **Reset.** Clears your changes on the tab or page and restores the most recently saved values.
- **Restart all servers on node.** Stops all application servers on the node and starts them again. The node agent for the node is also stopped and restarted.
- **Resume.** In the context of JCA lifecycle management, restarts the activity of a data source or a connection factory that was paused by a previous JCA lifecycle management operation.
- **Retrieve new.** Retrieves a new record.
- **Rollout update.** Sequentially updates an application that is installed on multiple cluster members across a cluster. After you update application files or a configuration, click **Rollout update** to install the configuration or the updated files for an application on all the cluster members of a cluster on which the application is installed. The Rollout update option applies the following steps to each cluster member in sequence:
 1. Saves an updated configuration.
 2. Stops the cluster member.
 3. Updates the application on the node by synchronizing the configuration.
 4. Restarts the cluster member.

Use **Rollout Update** if the application is deployed on one or more clusters spread across multiple nodes. This action reduces the amount of time that any single cluster member is unavailable to serve requests to the smallest interval possible. For a single node environment, use **Update** and then save and synchronize the node instead.

- **Save.** Saves the changes in your local configuration to the master configuration.
- **Select.** For resource analysis, lets you select a scope in which to monitor resources.
- **Set.** Saves your changes to settings in a dialog.
- **Settings.** Displays a dialog for editing servlet-related resource settings.
- **Settings in use.** Displays a dialog showing the settings in use.

- **Show groups.** Displays a collection of high availability groups, based on the match set.
- **Show servers.** Displays a collection of servers that are contained in the high availability groups that match the match set.
- **Start.** In the context of application servers, starts selected application servers. In the context of data collection, starts collecting data for the tables on this tabbed page.
- **Stop.** In the context of server components such as application servers, stops the selected server components. In the context of a data collection, stops collecting data for the tables on a tabbed page. In the context of nodes, stops servers on the selected nodes. In the context of deployment managers, stops the deployment manager server.
- **Synchronize.** Synchronizes the user configuration immediately. Click Synchronize on the Nodes page if automatic configuration synchronization is disabled, or if the synchronization interval is set to a long time, and a configuration change is made to the cell repository that needs replicating to that node. A node synchronization operation is performed using the normal synchronization optimization algorithm. This operation is fast, but might not fix problems from manual file edits that occur on the node. It is possible for the node and cell configuration to be out of synchronization after this operation is performed. If problems persist, use Full Resynchronize.
- **Terminate.** Deletes the Application Server process or another process that cannot be stopped by the **Stop** or **Immediate Stop** commands. Some application clients can receive exceptions. Always attempt an immediate stop before using this option.
- **Test connection.** After you define and save a data source, you can select this option to ensure that the parameters in the data source definition are correct. On the Collection panel, you can select multiple data sources and test them simultaneously.
- **Uninstall.** Deletes a deployed application from the WebSphere Application Server configuration repository. Also deletes application binary files from the file system.
- **Update.** For applications, replaces an application that is deployed on a server with an updated application. As part of the updating, you might need to complete steps on the Preparing for application installation and Update application pages.
For clusters, changes the configured weight or runtime weight assigned to a cluster member.
- **Update resource list.** Updates the data on a table. Discovers and adds new instances to the table.
- **Use cell CSI.** Enables Object Management Group (OMG) Common Secure Interoperability (CSI) protocol.
- **Use cell SAS.** Enables IBM Secure Authentication Service (SAS).
- **Use cell Security.** Enables cell security.
- **Verify tables.** Validates the mapping between the table names, scheduler resource, and data sources.
- **View.** Opens a dialog on a file.

Administrative console page features

This topic provides information about the basic elements of an administrative console page, such as the various tabs.

Administrative console pages are arranged in a few basic patterns. Understanding their layout and behavior can help you use them more easily.

Collection pages

Use collection pages to manage a collection of existing administrative objects. A collection page typically contains one or more of the following elements:

Scope Scope is described in Administrative console scope settings.

Preferences

Preferences are described in Administrative console preference settings.

Table of existing objects

The table displays existing administrative objects of the type specified by the collection page. The table columns summarize the values of the key settings for these objects. If no objects exist yet, an empty table is displayed. Use the available options to create a new object.

Buttons for performing actions

The available actions are described on the Administrative console buttons help panel. In most cases, you need to select one or more of the objects in the table, then click an action. The action is applied to the selected objects.

Sort toggle buttons

The column headings in the table are followed by icons for sort ascending (^) and sort descending (v). By default, items such as names are sorted in descending order (alphabetically). To enable another sorting order, click the icons for the column that you want to sort.

Detail pages

Use detail pages to configure specific administrative objects, such as an application server. A detail page typically contains one or more of the following elements:

Configuration tabbed page

This tabbed page is for modifying the configuration of an administrative object. Each configuration page has a set of general properties that is specific to the administrative object. Other sets of properties display on the page, but vary depending on the administrative object.

Runtime tabbed page

This tabbed page displays the configuration that is currently in use for the administrative object. The object is read-only in most cases. Some detail pages do not have runtime tabs.

If you can edit runtime properties, these properties directly affect the current runtime environment, but are not preserved when that environment is stopped.

Local Topology tabbed page

This tabbed page displays the topology that is currently in use for the administrative object. View the topology by expanding and collapsing the different levels of the topology. Some detail pages do not have local topology tabs.

Buttons for performing actions

Buttons to perform specific actions display on the configuration tabbed page and the Runtime tabbed page. The displayed buttons vary based on the administrative object. The available buttons are described on the Administrative console buttons help panel.

Wizard pages

Use wizard pages to complete a configuration process comprised of several steps. Be aware that wizards show or hide certain steps depending on the characteristics of the specific object that you are configuring.

Console layout

This topic describes the layout of the user interface for Integrated Solutions Console.

See Navigating the console for instructions on how to use the console controls.

Banner

Displays a common image across all Integrated Solutions Console installations. The banner includes a greeting to the user who is logged in and links to log out of the console and to open console help.

Navigation tree

Lists the tasks available in the console. Tasks are grouped into organizational nodes that represent categories of tasks, for example, Servers, or Applications. The organizational nodes can be nested in multiple levels.

The tasks shown are only those for which the user has access. When you click a task in the navigation, a page is displayed in the work area containing one or more modules for completing the task. Use the View selection list at the top of the navigation area to modify the list of tasks according to your preferences. You can organize the tasks as follows:

All tasks

This shows all tasks in the console. Tasks are grouped into organizational nodes, for example, Guided activities, Servers, or Applications.

My tasks

This shows only the tasks that you have added to the view. This list is initially empty, but provides a link to the **My Tasks** module. Use **My Tasks** to add and remove from the My Tasks list in the navigation.

Product selection

Selecting a product name shows only the tasks for that particular product, for example, WebSphere Application Server.

Work area

When you launch a page, the content of the page is displayed in the work area. If you have not launched any pages, the Welcome page is displayed in the work area. A page contains one or more console modules that are used to perform operations. Each console module has its own navigation controls. Some pages include a control to close the page and return to the Welcome page.

Console navigation

This topic describes how to navigate pages and tasks in the Integrated Solutions Console.

- Launching pages from the navigation tree
- Filtering tasks in the navigation
- Using the title bar controls
- Accessing help
- Using the console help controls



Before reading this section, read the terms and information in Console layout.

Launching pages from the navigation tree

The console navigation provides a hierarchical view of all of the *tasks* available in the console. A task is a page in the work area consisting of one or more console modules. All of the modules on the page are provided to start and complete the task. To open a task, simply click the task name in the navigation. The task is opened in a new page in the work area.

The following table describes the controls for the console navigation tree and entries in the tree.

Table 1. Console navigation tree control functions. The following table describes the controls for the console navigation tree and entries in the tree.









Icon	Function
	Represents an organizational node in the navigation tree that contains pages or other navigation nodes. Click the icon to expand the node.
	Closes an organizational node.

Filtering tasks in the navigation

When you first access the console, all tasks to which you have access are displayed in the navigation. Use the view menu at the top of the navigation to filter the list of tasks by product. Or, you can create a customized list by selecting My tasks from the **View** menu. For instructions on creating and managing your custom list of tasks, see My tasks.

Using the title bar controls

Each page contains one or more web applications or *console modules*. A console module enables you to perform an operation, such as displaying a list or stopping a managed system. The title and the controls for the module are displayed on the title bar. Depending on the functions supported by the module, the following icons might be displayed on the title bar:

- The  icon is displayed if the module allows you to edit settings for the portlet. For example, a module that retrieves performance data could permit you to specify the server to be analyzed. When you click the icon, an edit screen is displayed. Click the  icon to return to the previous screen.
- The  icon allows you to return to the previous screen.
- The  icon is displayed if help is available for the module. When you click the icon, the help is displayed in a separate browser window.
- The  icon allows you to minimize the module view. When you click the icon, only the title bar is visible on the page. Click the  icon to return to the maximize state.
- The  icon allows you to maximize a module view. When you click the icon, the full portlet view is visible on the page. Click the  icon to return to the minimize state.

In addition to the controls on the title bar, a module can include controls for other actions, such as a button to submit input. Some modules have controls that launch other modules. If a module launches another module, the newly launched module is displayed on a new page.


Accessing help

Help is available for the entire console or for a specific module in the console.

To access console help, perform the following steps:

1. Click **Help** on the console toolbar. The Help is displayed in a separate browser window.
2. In the help navigation tree, click the help set you want to view. For example, click **Console help** to view topics that provide helpful information for new console users. Use the console help controls as needed.

To access help for a module on a page, perform the following steps:

1. On the title bar for the module, click the  icon. That icon is displayed only if help is available for the module. The help is displayed in a separate browser window.
2. Close the help window when you are finished viewing it.

Using the console help controls

Table 2. Console help control functions. The following table describes the console help control functions.











Icon	Function
	Use these controls to navigate the list of pages you have viewed. Click  to return to the previous help topic that was displayed. Click  to move forward in the history list.
	Click either of these icons to synchronize the navigation tree with the current topic. The current topic will be highlighted in the navigation tree. This function is useful if you followed links from one help topic to other topics and you want to determine where the current topic is listed in the help navigation tree.
	Permits you to add the current page to your browser favorites list or bookmarks.
	Displays a window for printing the help topic that is displayed.

Table 2. Console help control functions (continued). The following table describes the console help control functions.

Icon	Function
	Maximizes the target view. This control is available for the Table of Contents view, the Search Results view, and the topic display area.
	Restores a maximized view to its normal size.
	Changes the view to the Search Results view. To search all of the help topics, type a word or words in the Search field. Enclose a phrase within double quotes. You can use Boolean operators (such as OR) in the search string. To limit the scope of the search, click Search scope . Click GO to start the search. A list of topics that contain the target strings are displayed in the results frame.
	Changes from the Search Results view to the Table of Contents view.

Administrative console browser support

Several web browsers are supported for use with Integrated Solutions Console.

The following web browsers are supported for use with Integrated Solutions Console:

- Firefox Versions 2.0, 3.0, 3.5, and 3.6
- Firefox Version 1.5, for AIX V6.1 only
- Microsoft Internet Explorer Versions 7.0 and 8.0

Note: Using the browser's back button with the console can produce unexpected results and is not supported. Use the controls and links provided in the console to navigate between pages and applications.

Console accessibility

There are many accessibility features built into Integrated Solutions Console.

- Accessibility features
- Navigating the console by using the keyboard
- Navigating help by using the keyboard

Accessibility features

The Integrated Solutions Console has the following accessibility features:

- The following features are for vision-impaired users:
 - Can be operated by using only the keyboard
 - Communicates all information independent of color
 - Supports the attachment of alternate output devices
 - Provides help information in an accessible format
- The following features are for users who have mobility impairments or limited use of their hands:
 - Allows the user to request more time to complete timed responses
 - Can be operated by using only the keyboard
 - Supports the attachment of alternative input and output devices
- The following features are for the deaf and hard of hearing users:
 - Supports alternatives to audio information
 - Supports adjustable volume control
- The console does not flash the screen at rates that could induce epileptic seizures.

The help system for Integrated Solutions Console has the following accessibility features:

- Uses the accessibility support enabled by the browser that is used to display the help
- Enables navigation by using the keyboard

Navigating the console by using the keyboard

To move through the controls on a particular page, use the Tab key.

To click a link or control on a page using the keyboard, navigate to the link or control and press Enter.

To change the navigation view using the keyboard, follow these steps.

1. Navigate to the **View** selection list using the Tab key.
2. Use the up and down arrows to change the value of the selection list.
3. Press Enter. The tasks displayed in the navigation are changed according to your selection.

Navigating help by using the keyboard

Use the following key combinations to navigate the help system by keyboard:

- To bring the Topic pane (the right hand side) into focus, press Alt+K, and then press Tab.
- In the Topic pane, to go to the next link, press Tab. To go to the previous link, press Shift+Tab.
- To go directly to the Search Results view in the left hand side, press Alt+R, and then press Enter or Up arrow to enter the view.
- To go directly to the Navigation (Table of Contents) view in the left hand side, press Alt+C, and then press Enter or Up arrow to enter the view.
- To navigate your browser history, press Alt+Left arrow to go back. If you have navigated back to a previously view page, you can use Alt+Right arrow to navigate forward again.
- To expand and collapse a node in the navigation tree, tab to the + or - image next to it to bring the image into focus, and then press the Right or Left arrows.
- To go to the next frame in the help system, press F6. To go to the previous frame in the help system, press Shift+F6.
- In the navigation, to move to the next topic node, press the Down arrow or Tab. To move to the previous topic node, press the Up arrow or Shift+Tab.
- To go to the next link, button, or topic node from inside a view, press Tab.
- To scroll all the way up or down in a frame, press Home or End, respectively.
- To print the active pane, press Ctrl+P.
- To move to the search entry field, press Alt+S.

Welcome

Display products that are installed that use the Integrated Solutions Console for administrative tasks.

The Welcome page displays the products that are installed that use the Integrated Solutions Console for administrative tasks. The page lists the product name and version number. If provided by the product, you can click the product name to display a page that provides more information about the product.

My tasks

Create and edit a task view in the console navigation.

Use My tasks to create and edit a list of tasks to view in the console navigation. A task includes a page that contains one or more web applications, or *console modules*, that are used to complete that task. When you first access the console, all tasks to which you have access are displayed in the navigation. My tasks is especially useful to customize the navigation to show only the tasks you use most often. After you customize your tasks, My Tasks is initially displayed each time you log in to the console.

Follow these general steps to customize your task list in the navigation.

1. Select My tasks from the **View** selection list in the navigation. If you have never used My tasks before, you must click **Add tasks** to open it.
2. Use the checkboxes to select and deselect tasks from the My tasks navigation.
3. To save your changes, click **Apply**.
4. To cancel your changes, click **Reset**.

After applying your selections, your customized task list is displayed in the navigation.

Use the following buttons to customize your task selections.

Apply Saves the current selections.

Reset Backtracks all changes to the selections that were set since the last time My tasks was applied. This is useful if you need to cancel your changes.

Select All
Checks every task.

Deselect All
Unchecks every task.

Expand All
Expands each node in the display and reveals all subtasks in the navigation.

Collapse All
Collapses each node in the display so that only the top level nodes are displayed.

Console identity

Use Console Identity to define a string (keyword, name, or phrase) to be displayed in the console banner. This string can be used to distinguish this console from other console instances.

To access this page, click **System administration > Console Identity** in the console navigation.

The console identity string is rendered differently in the banner and browser title bar.

- **Console identity in the banner**

The console identity string is displayed in the console banner after the greeting to the console user, separated by a dash. For example, if the identity string is set to *Jupiter* and the user is logged in as *consoleadmin*, the banner displays the following greeting.

Welcome consoleadmin - Jupiter

On the login page, the console identity is displayed without the greeting.

- **Console identity in the browser title bar**

The console identity string is displayed in the title bar after the console brand name, separated by a space. For example, if the identity string is set to *Saturn*, the title bar displays the following information.

Integrated Solutions Console Saturn

Note: After saving these changes, the console identity settings are applied across the console. For an administrative agent configuration, this means that the changes are applied to the administrative agent and all of its registered application servers, regardless of where the changes were actually saved.

Console identity:

Specifies whether a custom string is used to identify this console.

none	Select this option if a custom string should not be displayed.
custom	Select this option to display a custom string in the banner and browser. When this option is selected, Custom identity string is enabled.

Custom identity string:

Specifies the text string that you want to display in the banner and browser title bar.

Truncate string at (Characters):

Specifies how many characters of the string should be displayed. Input for this field must be an integer. Minimum value is 4, maximum value is 99. When the string is truncated, 3 of the characters are used for an ellipsis (...) to indicate that it has been truncated.

Console identity preview:

Click **Preview** to see how the custom string will be displayed in the banner.

Console identity string

You can define a console identity string to be displayed in the console banner and in the browser's title bar. This string can be fixed or you can use a variable that is resolved at run time with a Java system property or environment variable.

The console identity string is rendered differently in the banner and browser title bar.

- **Console identity in the banner**

The console identity string is displayed in the console banner after the greeting to the console user, separated by a dash. For example, if the identity string is set to `Jupiter` and the user is logged in as `consoleadmin`, the banner displays the following greeting.

Welcome consoleadmin - Jupiter

On the login page, the console identity is displayed without the greeting.

- **Console identity in the browser title bar**

The console identity string is displayed in the title bar after the console brand name, separated by a space. For example, if the identity string is set to `Saturn`, the title bar displays the following information.

Integrated Solutions Console Saturn

After installation, there are two ways provided to customize the identity string.

- Use Console Identity in the administrative console. To open, log in to the administrative console and click **System administration > Console Identity**.
- Set the value for the string in an XML file, `consoleProperties.xml`, which is located in `app_server_root/profiles/profile_name/config/cell/cell_name/applications/isclite.ear/deployments/isclite/isclite.war/WEB-INF`. You should keep a backup copy of the most recent working `consoleProperties.xml` before making changes to this file.

The keys that are used to set the console identity string to be displayed in the banner and title bar are defined in an XML `<console-property/>` element. The required attributes **id** and **value** for this element determine the key name and value pairs. The following initial settings are provided in `consoleProperties.xml`.

(Some keys are split on multiple lines for printing purposes.)

Table 3. Console identity key values. The following table describes the console identity key values.

Key (id)	Initial setting (value)	Description
ISC.CONSOLE.ID	ISC.CONSOLE.ID.CUSTOM	Specifies the ID attribute of another <code><console-property/></code> element that contains the actual console identity string
ISC.CONSOLE.ID.MAXLEN	27	Specifies the maximum length of the resolved console identity string. The string will be truncated to this length if necessary.
ISC.CONSOLE.ID.CUSTOM	<i>empty string</i>	A custom string to be displayed in the banner and title bar.
ISC.CONSOLE.ID.DEPLOYER. <i>unique_suffix</i>	<i>None</i>	Specifies a custom string that can be selected using the Console Identity application under System Administration. By providing multiple <code><console-property/></code> entries with the <code>ISC.CONSOLE.ID.DEPLOYER. <i>unique_suffix</i></code> id attribute, you can provide multiple identity strings for console users to select from. Each entry of this type should have a distinct <i>unique_suffix</i> from the other entries. If multiple entries are provided with the same <i>unique_suffix</i> , then only the first value with this suffix is displayed in the selection list for the console user.

Examples:

- In the following example, the deployer wants to create custom strings that indicate different departments served by their respective host consoles. Console users will be able to select which identity is displayed in each department's console interface. To achieve this goal, the deployer provides the following custom keys in the installation package. Out of the box, the blank value for `ISC.CONSOLE.ID.CUSTOM` is used.

```
<?xml version="1.0" encoding="UTF-8"?>
<consoleproperties:ibm-portal-consoleproperties xmlns:consoleproperties=
  "http://www.ibm.com/websphere/appserver/schemas/6.0/ibm-portal-consoleproperties.xsd"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation=
    "http://www.ibm.com/websphere/appserver/schemas/6.0/ibm-portal-consoleproperties.xsd
      ibm-portal-consoleproperties.xsd">
```

```

<consoleproperties:console-property id="ISC.CONSOLE.ID"
                                   value="ISC.CONSOLE.ID.CUSTOM"/>
<consoleproperties:console-property id="ISC.CONSOLE.ID.MAXLEN" value="27"/>
<consoleproperties:console-property id="ISC.CONSOLE.ID.CUSTOM" value=""/>
<consoleproperties:console-property
    id="ISC.CONSOLE.ID.DEPLOYER.sales" value="Sales & Marketing"/>
<consoleproperties:console-property
    id="ISC.CONSOLE.ID.DEPLOYER.finance" value="Finance"/>
<consoleproperties:console-property
    id="ISC.CONSOLE.ID.DEPLOYER.research" value="Research & Development"/>
<consoleproperties:console-property
    id="ISC.CONSOLE.ID.DEPLOYER.dist" value="Distribution"/>
</consoleproperties:ibm-portal-consoleproperties>

```

- The administrator wants to append information about the host console to the banner and title bar. This can be accomplished by updating the ISC.CONSOLE.ID.CUSTOM key to include the PROCESSOR_ARCHITECTURE environment variable (which works on Windows machines) along with the **os.name** Java system property. The other keys are left unchanged.

```

<?xml version="1.0" encoding="UTF-8"?>
<consoleproperties:ibm-portal-consoleproperties
    xmlns:consoleproperties=
        "http://www.ibm.com/websphere/appserver/schemas/6.0/ibm-portal-consoleproperties.xsd"
    xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
    xsi:schemaLocation=
        "http://www.ibm.com/websphere/appserver/schemas/6.0/ibm-portal-consoleproperties.xsd
        ibm-portal-consoleproperties.xsd">
    <consoleproperties:console-property id="ISC.CONSOLE.ID"
                                       value="ISC.CONSOLE.ID.CUSTOM"/>
    <consoleproperties:console-property id="ISC.CONSOLE.ID.MAXLEN" value="27"/>
    <consoleproperties:console-property id="ISC.CONSOLE.ID.CUSTOM"
                                       value="Platform: %PROCESSOR_ARCHITECTURE% OS: %os.name%"/>
</consoleproperties:ibm-portal-consoleproperties>

```

Since the ISC.CONSOLE.ID.MAXLEN key indicates a maximum display of 27 characters, the resulting string, which is resolved at run time, is rendered in the banner as follows for a Windows 2003 Server environment.

Platform: x86 OS: Window...

Note: After the console identity string is saved, the settings are applied across the console. For an administrative agent configuration, this means that the changes are applied to the administrative agent and all of its registered application servers, regardless of where the changes were actually saved.

Administrative console: Resources for learning

Use the following links to find relevant supplemental information about the IBM WebSphere Application Server 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 IBM Redbooks that supplement the broad coverage of the release documentation with in-depth examinations of particular product areas.

View links to additional information:

Administration

- IBM WebSphere Application Server Redbooks

This site contains a listing of all WebSphere Application Server Redbooks.

- IBM WebSphere developerWorks

This site is the home of technical information for developers working with WebSphere products. You can download WebSphere software, take a fast path to developerWorks zones, such as VisualAge® Java or WebSphere Application Server, learn about WebSphere products through a newcomers page, tutorials, technology previews, training, and Redbooks, get answers to questions about WebSphere products, and join the WebSphere community, where you can keep up with the latest developments and technical papers.

- WebSphere Application Server Support page

Take advantage of the Web-based Support and Service resources from IBM to quickly find answers to your technical questions. You can easily access this extensive Web-based support through the IBM Software Support portal at web address <http://www.ibm.com/software/support/> and search by product category, or by product name. For example, if you are experiencing problems specific to WebSphere Application Server, click **WebSphere Application Server** in the product list. The WebSphere Application Server Support page appears.

Installing and uninstalling the administrative console

You can install the administrative console during profile creation or after you create a profile. You can uninstall any administrative console that you install. To install an administrative console after profile creation, or to uninstall the administrative console, use the **wsadmin** command. This topic discusses how to use the **wsadmin** command to install and uninstall the administrative console.

Before you begin

If you install the administrative console through the **wsadmin** command, a profile that does not have an administrative console installed must exist.

About this task

Run the `deployConsole` script on the **wsadmin** command whenever you want to uninstall the administrative console, or whenever you want to install the administrative console to a profile that does not have an administrative console installed.

You can run the script in either connected or disconnected mode.

Unfederated application servers, administrative agents, deployment managers, and job managers can have their own administrative consoles. The steps in this task apply to these consoles.

The usual security restrictions for the **wsadmin** command apply to this script. In connected mode, the user must authenticate if security is enabled.

To install or uninstall the deployment manager administrative console, the **wsadmin** command attempts to remotely connect to the deployment manager. However, various situations can keep the **wsadmin** command from connecting to the deployment manager, and an error message results. The command cannot connect if:

- The deployment manager is not running.
- The deployment manager is running, but the script was run somewhere other than on the deployment manager. You are most likely running the script on a federated node.

The `deployConsole.py` script is located in the `profile_root/bin` directory.

Procedure

- To install the administrative console, issue the following command:
`wsadmin.sh -f deployConsole.py install`
- To uninstall the administrative console, issue the following command:
`wsadmin.sh -f deployConsole.py remove`

Results

The administrative console is installed or uninstalled, depending on whether you specified the install or remove option.

Starting and logging off the administrative console

This topic describes how to set up the administrative console environment, to access the administrative console, and to log out of the administrative console.

Before you begin

To access the administrative console, you must first install WebSphere Application Server and the administrative console.

About this task

Note: This topic references one or more of the application server log files. Beginning in WebSphere Application Server Version 8.0 you can configure the server to use the High Performance Extensible Logging (HPEL) log and trace infrastructure instead of using `SystemOut.log`, `SystemErr.log`, `trace.log`, and `activity.log` files or native z/OS logging facilities. If you are using HPEL, you can access all of your log and trace information using the LogViewer command-line tool from your server profile bin directory. See the information about using HPEL to troubleshoot applications for more information on using HPEL.

The administrative console application is installed during the initial installation process.

You start the administrative console, access the console through a web browser, and then log into the administrative console. After you finish working in the administrative console, save your work and log out.

Procedure

1. Start the desired administrative console by starting the server process that runs the console application.
You can start an unfederated application server, an administrative agent server, a deployment manager server, or a job manager server.
Check the `SystemOut.log` file of the server that runs the console application to verify that the console application starts successfully. If the console application starts successfully, you see the `WSVR0221I: Application started: isclite` message.
2. Access the administrative console.
 - a. Enable cookies in the web browser that you use to access the administrative console.
 - b. Enable JavaScript.
Enablement of JavaScript is required. You must enable JavaScript so that all the features of the administrative console are available.
 - c. In the same web browser, type `http://your_fully_qualified_server_name:port_number/ibm/console`.
The `your_fully_qualified_server_name` string is the fully qualified host name for the machine that contains the administrative console. The administrative console can be for an unfederated

application server, a deployment manager, an administrative agent, or a job manager. The port number for the *port_number* string is 9060 by default for the unfederated application server, the deployment manager, and the administrative agent. For the job manager, the port number for the *port_number* string is 9960 by default.

If you cannot start the administrative console because the console port conflicts with an application that is already running on the machine, do one of the following actions:

- Change the port number and propagate the number to the appropriate files:
 - 1) Change all the occurrences of the console port to a new port number. Make the port changes in the *installation root/profiles/profile name/config/cells/cell_name/nodes/node_name/servers/server_name/serverindex.xml* file and the *installation root/profiles/profile name/config/cells/cell_name/virtualhosts.xml* files.
 - 2) Run the *./wsc2n.sh* script from the *installation root/WebSphere/AppServer/bin* directory. The *./wsc2n.sh* script generates the *was.env* file, the *control.jvm.options* file, the *servant.jvm.options* file, and the *adjunct.jvm.options* file for each server and the *was.env* file for the location service daemon. These generated files will contain the updated administrative console port number.
- Shut down the other application that uses the conflicting port before starting the WebSphere Application Server product.

For a listing of supported web browsers, see WebSphere Application Server system requirements at

<http://www.ibm.com/support/docview.wss?rs=180&uid=swg27006921>

The web address displays on two lines for printing purposes. Enter the web address on one line in your browser.

- d. Wait for the administrative console to load into the browser.

A login page displays after the administrative console starts.
3. Log into the administrative console.

The administrative console can be for an unfederated application server, a deployment manager, an administrative agent, or a job manager.

- a. If you are logging into the administrative console for the administrative agent, and you have registered at least one node with the administrative agent, select the node to administer, and click **Continue**.

The node can be the administrative agent node or a node for one of the application servers registered to the administrative agent. After you select a node, the login procedure is the same as that for the other server types, and for administrative agents with no nodes registered.

- b. Enter your user name or user ID.

The user ID lasts only for the duration of the session for which it is used to log in.

Changes made to server configurations are saved to the user ID. Server configurations also are saved to the user ID if a session timeout occurs.

See the *Administering applications and their environment* PDF for more information on server configurations.

If you enter an ID that is already in use and in session, you are prompted to do one of the following actions:

- Log out the other user with the same user ID. You can recover changes made during the other user's session.
 - Return to the login page and enter a different user ID.
- c. If the console is secure, you must also enter a password for the user name. The console is secure if someone has taken the following actions for the console:
 - Specified security user IDs and passwords
 - Enabled global security

See the *Securing applications and their environment* PDF for more information.

d. Click **OK**.

4. Log off the administrative console. Click **System administration > Save changes to master repository > Save** to save work. Then click **Logout** to exit the console.

If you close the browser before saving your work, you can recover any unsaved changes the next time that you log in under the same user ID.

Results

You have set up the administrative console environment, accessed the administrative console, and logged out of the administrative console.

What to do next

Use the administrative console to manage the product.

Logging in

Enter your user ID and password to access the console.

To access the console, enter your **User ID** and **Password** and then click **Log in**. The password is required only if security is enabled. In environments that use the administrative agent to administer multiple application server nodes, select whether to log in to the administrative agent or one of its registered profiles.

After you are logged in, be sure to use the **Logout** link in the console toolbar when you are finished using the console and to prevent unauthorized access. If there is no activity during this login session for an extended period of time, the session expires and you must login again to access the console. The administrator can change the session timeout. The default is set to 30 minutes.

If the user ID that you provide is already logged in at a different location, you are prompted to choose between logging out from the other location or returning to the login page. If you log out the user from the other location, you might be prompted to recover unsaved changes made by that user.

If you have one or more different stand-alone servers running on the same machine and wish to administer them concurrently from the same or a different machine then you must:

1. Ensure that each server uses a unique value for its admin console port.
2. Run a separate web browser process for each admin console that you wish to access concurrently.

Save changes to the master repository

Use this topic to update the master repository with your administrative console changes, to discard your administrative console changes and continue working with the master repository, or to continue working with your administrative console changes that are not saved to the master repository.

Until you save changes to the master repository, the administrative console uses a local workspace to track your changes.

Total changed documents: Specifies the total number of documents that you changed for your session, but that are not saved to the master repository. By clicking the +/- toggle key, you can see additional information about the changed documents:

- **Changed items**

When you change your local configuration, each path and configuration file that you can apply the update to in the master repository is displayed in the list.

- **Status**

The status can contain the following options:

- **Added:** If you save your changes to the master repository, a new configuration file is created on the indicated path.
- **Updated:** If you save your changes to the master repository, an existing configuration file is updated on the indicated path.
- **Deleted:** If you save your changes to the master repository, an existing configuration file is deleted on the indicated path.

Synchronize changes with nodes: Specifies whether you want to force node synchronization at the time that you save your changes to the master repository, rather than when node synchronization normally occurs.

Save conflict: Specifies that another user changed some configuration information since you began making changes. You can either click **Save** to overwrite the other user information, or **Discard** to discard your changes and keep the changes that the other user made.

Specifying console preferences

Use this topic to customize how much data displays on an administrative console panel.

About this task

Throughout the administrative console are pages that have Preferences fields, Scope fields, and Filter radio buttons. By selecting these fields and radio buttons you can customize how much data is shown.

For example, examine the Preferences field for the Administrative authorization groups page:

Procedure

1. Go to the navigation tree of the administrative console and click **Security > Administrative authorization groups**.
2. Expand **Preferences**.
3. For the **Maximum rows** field, specify the maximum number of rows to display when the collection is large. The default is 20. Rows that exceed the maximum number display on subsequent pages.
4. Select **Retain filter criteria** if you want to retain the last filter criteria that is entered in the filter function. When you return to the Applications page, the page initially uses the retained filter criteria to display the collection of applications in the table following the preferences. Otherwise, clear **Retain filter criteria** and the last filter criteria is not retained.
5. Click **Apply** to apply your selections or click **Reset** to return to the default values. The default is not to enable (not have a check mark beside) **Retain filter criteria**.

Results

Other pages have similar fields and radio buttons that you can use to specify console preferences. While Preferences fields, Scope fields, and Filter buttons control how much data is shown in the console, the **Preferences** option controls general behavior of the console. Click **System Administration > Console preferences** to view the Preferences page.

Console preferences settings

Use the Console Preferences page to specify how you want features of the administrative console workspace to behave.

To view this administrative console page, click **System administration > Console preferences**.

Turn on workspace automatic refresh:

Specifies whether you want the administrative console workspace to refresh automatically after the administrative configuration changes.

The default is for the workspace to refresh automatically. If you delete a WebSphere variable, for example, the WebSphere variables page refreshes automatically and shows the updated list of WebSphere variables in the WebSphere variables collection.

Specifying that the workspace not refresh automatically means that you must access a page again by clicking the console navigation tree or links on collection pages to see the changes that are made to the administrative configuration.

Default true (selected)

No confirmation on workspace discard:

Specifies whether the confirmation dialog is displayed after a request is received to discard the workspace. The default is to display confirmation dialogs.

Default false (cleared)

Use default scope:

Specifies whether the default scope is the administrative console node.

All scopes is the default unless you enable the Use default scope setting to make the administrative console node the default. Whatever the default is the first time that you view a console panel that has scope settings, that is the default for the panel on subsequent visits that you make to the panel. The default for the panel does not change even if you modify the Use default scope setting.

Default false (cleared)

Show the help portlet: Specifies whether the help portlet on the right of the console displays.

Default true (selected)

Enable command assistance notifications: Specifies whether to send Java Management Extensions (JMX) notifications that contain command assistance data from the administrative console. Enablement of the notifications allows integration with product tools such as the Toolkit Jython editor for WebSphere Application Server. Enablement of this option is recommended for non-production environments only.

Default false (cleared)

Log command assistance commands: Specifies whether to log all the command assistance wsadmin data to a file. This file is saved to `${LOG_ROOT}/server/commandAssistanceJythonCommands_ user name.log`:

- *server* is the server process where the console runs, such as dmgr, server1, adminagent, or jobmgr.
- *user name* is the administrative console user name.
- When you manage a profile using an administrative agent, the command assistance log is put in the location of the profile that the administrative agent is managing. The `${LOG_ROOT}` variable defines the profile location.

Occasionally clean out the file to manage its growth.

Default false (cleared)

Synchronize changes with nodes: Specifies whether to synchronize changes that are saved to the deployment manager profile with all the nodes that are running.

Default false (cleared)

Bidirectional support options: Specifies bidirectional (Bidi) text preferences for the administrative console.

Default false (cleared)

Bidirectional support options

Use the Bidirectional support options page to specify bidirectional (Bidi) text preferences for the administrative console.

Bidirectional support means that text is supported going in both directions for different types of alphabets. WebSphere Application Server presents Bidi text using left-to-right (LTR) orientation in most languages. However, when Bidi text is used as part of mixed Bidi and Latin text, for example, the preferred orientation is right-to-left (RTL).

For text entry fields, this means that the cursor is placed at the right side and moves to the left as characters are typed in an RTL alphabet, such as Hebrew or Arabic.

Bidi support helps to maintain visual structure in complex fields, such as file paths, emails, URLs and Xpaths.

If you select **Enable bidirectional support for all users** under Global Preferences, any change made affects all users. Note that if you select this option that **Enable bidirectional support for this user** under Current[®] User Preferences is automatically selected also. However, if you only select **Enable bidirectional support for this user**, any changes made only affect the user that is currently logged in.

To view this administrative console page, click **Environment > Console preferences > Bidirectional support options**.

Enable bidirectional support for all users:

Select this choice to enable bidirectional support for all users.

Default false

Default text direction for all users:

Select one of the options on this menu to indicate which direction the text should go for all users.

The default text direction is left to right. This is the natural base text direction for most languages, including European, Asian and Indic languages

If you select right to left, this is the natural base text direction for languages such as Arabic, Hebrew, Urdu and Farsi.

If you select contextual, the base text direction is set according to the first strong character. This is an appropriate choice when at the time of setting the base direction, the main language of the string is not yet known.

Default LTR

Enable bidirectional support for this user:

Select one of the options in this menu to indicate which direction the text should go for only the user currently logged in.

Default false

Text direction for this user: The default text direction is left to right. This is the natural base text direction for most languages, including European, Asian and Indic languages

If you select right to left, this is the natural base text direction for languages such as Arabic, Hebrew, Urdu and Farsi.

If you select contextual, the base text direction is set according to the first strong character. This is an appropriate choice when at the time of setting the base direction, the main language of the string is not yet known.

Default LTR

Administrative console preference settings

Use the preference settings to specify how you want information to display on an administrative console panel. The preference settings vary from one administrative console panel to another.

Maximum rows: Indicates the maximum number of rows to display per page when the collection is large.

Filter history: Indicates whether to use the same filter criteria to display this page the next time that you visit.

Select the **Retain filter criteria** check box to retain the last filter criteria entered. When you return to the page, retained filter criteria control the application collection that is displayed in the table.

Show resources in the scope hierarchy: Select the check box if you want to display the resources in the hierarchy for a particular scope.

The hierarchy is:

- cell > node > server
- cell > cluster

For example, if you select a node scope, all node scope resources and all cell scope resources display for the node.

This preference is available for resource factory panels only.

Show built-in resources: Select the check box if you want to display resources that are pre-defined to support certain internal components of the product. For example, the product includes built-in configurations of a Cloudscape JDBC provider and a data source to support the Universal Description, Discovery and Integration Protocol (UDDI) registry for web services.

Show confirmation for stop command: Select the check box if you want a confirmation that the **stop** command is successful.

Show confirmation for immediate stop command: Select the check box if you want a confirmation that the **immediate stop** command is successful.

Display inherit policy set attachments confirmation: Select the check box if you want to enable the inherit policy set attachments confirmation.

Show confirmation for terminate command: Select the check box if you want a confirmation that the **terminate** command is successful.

Show resources at one authorizing group level only: Specifies the authorization group level used to filter the resources in the table. Only those roles that apply to your ID can display in the table. Valid values are All Roles, Administrator, Deployer (for application collection panels only), Operator, Configurator, and Monitor. If All Roles is selected, then all the resources that you are authorized to view are displayed in the table grouped by role. Otherwise, the resources for the role selected display in the table.

Hide system queues: Select the box to hide messaging engine system queue points.

Include cluster members in the collection: Select the check box if you want the collection to include application servers that belong to a server cluster.

Show confirmation for update runtime command: Select the check box to enable the confirmation panel for the update runtime command button.

Show items at the following authorization group level: Select from the list the authorization group level that will be used to filter the items in the table.

Show all data source properties: By default, this panel does not list the custom properties that are configurable by different administrative console pages. Select the check box to show all data source custom properties, including properties required by the data source and configuration properties. Selecting this option does not affect what is displayed on other panels.

Administrative console scope settings

Use this page to specify the level at which a resource is visible on the administrative console panel. By changing the value for Scope, you see only the resources that are defined at that scope. The contents of the collection table might change. For WebSphere Application Server, Network Deployment, a resource can be visible in the administrative console collection table at the cell, node, cluster, or server scope.

For WebSphere Application Server, Network Deployment, if you select **Show scope selection drop-down list with the all scopes option**, the console displays a drop-down list of all the scopes available. To change the scope, select any item from the drop-down list. The drop-down list contains only those scopes that you have the authority to view.

For WebSphere Application Server, Network Deployment, if you clear **Show scope selection drop-down list with the all scopes option**, you do not receive the drop-down list. Instead, click **Browse** next to a field to see choices for changing the scope of the field. If a field is read-only, you cannot change the scope. For example, if only one server exists, you cannot switch the scope to a different server.

The value for the **Show scope selection drop-down list with the all scopes option** checkbox is saved on a per-user and per-resource basis. For example, user 1 might select the checkbox while user 2 clears the checkbox. The values for the checkbox are saved for each user separately. Different resources, for example Java Database Connectivity (JDBC) Providers and data sources, can have different values for the checkbox.

All scopes is the default unless you enable the Use default scope setting on the Console preferences panel to make the administrative console node the default. Whatever the default is the first time that you view a console panel that has scope settings, that is the default for the panel on subsequent visits that you make to the panel. The default for the panel does not change even if you modify the Use default scope setting.

You cannot select All scopes to create a new resource. You must select one of the available scopes from the drop down list to create a new resource.

You always create resources at the current scope that is selected in the administrative console panel, even though the resources might be visible at more than one scope.

Resources such as Java Database Connectivity (JDBC) providers, namespace bindings, or shared libraries can be defined at multiple scopes. Resources that are defined at more specific scopes override duplicate resources that are defined at more general scopes:

- The application scope has precedence over all the scopes.
- For WebSphere Application Server, Network Deployment, the server scope has precedence over the node, cell, and cluster scopes.
- For WebSphere Application Server, Network Deployment, the cluster scope has precedence over the node and cell scopes.
- The node scope has precedence over the cell scope.

Despite the scope of a defined resource, the resource properties apply at an individual server level only. For example, if you define the scope of a data source at the cell level, all the 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 define the maximum connections as 10, then each server in that cell can have 10 connections.

The cell scope is the most general scope and does not override any other scope. The recommendation is that you generally specify a more specific scope than the cell scope. When you define a resource at a more specific scope, you provide greater isolation for the resource. When you define a resource at a more general scope, you provide less isolation. Greater exposure to cross-application conflicts occur for a resource that you define at a more general scope.

Cell Limits the visibility to all servers on the named cell. The resource factories within the cell scope are:

- Defined for all servers within this cell
- For WebSphere Application Server, Network Deployment, overridden by any resource factories that are defined within application, server, cluster, and node scopes that are in this cell and have the same Java Naming and Directory Interface (JNDI) name

The resource providers that are required by the resource factories must be installed on every node within the cell before applications can bind or use them.

Cluster

Clusters are used in WebSphere Application Server, Network Deployment. This scope limits the visibility to all the servers on the named cluster. All cluster members must at least be at Version 6 to use cluster scope for the cluster. The resource factories that are defined within the cluster scope:

- Are available for all the members of this cluster to use
- Override any resource factories that have the same JNDI name that is defined within the cell scope

The resource factories that are defined within the cell scope are available for this cluster to use, in addition to the resource factories, that are defined within this cluster scope.

Node Limits the visibility to all the servers on the named node. The node scope is the default scope for most resource types. The resource factories that are defined within the node scope:

- Are available for servers on this node to use
- Override any resource factories that have the same JNDI name defined within the cell scope

The resource factories that are defined within the cell scope are available for servers on this node to use, in addition to the resource factories that are defined within this node scope.

Server

Limits the visibility to the named server. The server scope is the most specific scope for defining resources. The resource factories that are defined within the server scope:

- Are available for applications that are deployed on this server
- Override any resource factories that have the same JNDI name defined within the node and cell scopes

The resource factories that are defined within the node and cell scopes are available for this server to use, in addition to the resource factories that are defined within this server scope.

Application

Limits the visibility to the named application. Application scope resources can be viewed and edited from the console, but not created. You can additionally use the Rational® Application Developer or the wsadmin tool to view or edit the application scope resource configuration. The resource factories that are defined within the application scope are available for this application to use only. The application scope overrides all other scopes.

You can view the application scope resources from the console by selecting **Applications** from the console navigation, and then navigating to the appropriate application. The application scope resources are unavailable from the Resources section of the console navigation.

You can configure namespace bindings and shared libraries under cell, node, and server scopes only. For WebSphere Application Server, Network Deployment, you can configure resources and the product variables under all five scopes.

Accessing help and product information from the administrative console

This topic describes how to use administrative console help and how to link to product documentation from the administrative console.

Before you begin

You must have a connection to the Internet to access information about WebSphere Application Server from the Welcome page of the administrative console.

About this task

All of the helps panels that you can access from the administrative console, you can access from the WebSphere Application Server Information Center. This topic describes how to access the help panels, the information center, and other product documentation from the administrative console.

Procedure

- Click **Welcome** on the administrative console navigation tree. In the workspace to the right of the navigation tree, click the link, which takes you to a page that has links to various documentation. The documentation that is linked includes the WebSphere Application Server Information Center, the WebSphere Application Server product information, and the WebSphere Application Server technical information on developerWorks.
- Access help in the following ways:
 - Click either of the following tabs of an online help page:
 - Click the **Help index** tab and select from the list of help panels to view administrative console help information.
 - Click the **Search** tab, provide search terms, and then click **Search**. Under Results, select a help panel that contains the search information.
 - In the help portal that is on the right side of the administrative console panel, do one or all of the following tasks:

- Click a field label or a list marker in the administrative console panel for the help to display under Field help. Alternatively, place the cursor over the field label or the list marker for the corresponding help to display at the cursor.

Attention: When you place the cursor over the field label or list marker, the help might be truncated in a Firefox browser. Click the field label or list marker so that the full help displays under Field help.

- Click the link under Page help to access the help panel for the administrative console panel. The help panel is the same help panel that displays when you click the ? icon.
- If Command assistance is listed, click the link under Command assistance to view wsadmin scripting commands for the last action run for this console panel.

What to do next

You can continue to access help information from the administrative console. Alternatively, you can access the help information from the WebSphere Application Server Information Center.

You can continue to access the WebSphere Application Server Information Center, the WebSphere Application Server product information, and the WebSphere Application Server technical information on developerWorks from the administrative console. Alternatively you can access the information from the IBM website.

Accessing command assistance from the administrative console

Using command assistance, you can view wsadmin scripting commands in the Jython language for the last action run in the administrative console. This topic discusses how to access command assistance from the administrative console.

Before you begin

You must have WebSphere Application Server and the administrative console running to access command assistance.

About this task

Use command assistance to see wsadmin scripting commands that correspond to actions in the administrative console. Seeing these commands might help you develop the commands necessary to administer WebSphere Application Server from the wsadmin utility.

If a command assistance link is listed in the help portlet, **wsadmin** commands exist for the last console action that you completed, and command assistance is available for that action.

When command assistance is unavailable in the help portlet: Some console actions do not have **wsadmin** commands directly associated with them. When the help portlet on the right side of the administrative console panel does not have a command assistance link in it, no command assistance data is available for the last console action.

Procedure

1. Click the link under **Command assistance** to view wsadmin scripting commands for the last action run for this administrative console panel.

After the Command assistance window opens, it refreshes automatically when new command assistance data is available.

Examples of actions include a click on a button or a click on a link in the navigation bar, a collection panel, or a detail panel. The editing of forms is not a user action.

The wsadmin scripting commands display in the Jython language in a secondary window that you can view by clicking on the Command assistance link in the help portlet.

If you perform an administrative console action after you launch the Command assistance window, whether or not the scripting commands display in the window depends on whether your browser supports Javascript. If your browser supports Javascript, the Command assistance window automatically refreshes the command list to reflect the most recent console action. If the browser does not support Javascript, click the link again under **Command assistance** in the help portal to refresh the command list.

2. To view the description of a specific **wsadmin** command, place your cursor over the command. Hover text is displayed.
3. Optionally, log the command assistance data to a file by selecting the **Log command assistance commands** setting on the Preferences page of the administrative console.
A timestamp and the breadcrumb trail of the panel that produced the command assistance data are provided with the wsadmin data.
4. Optionally, allow command assistance to emit Java Management Extensions (JMX) notifications by selecting the **Enable command assistance notifications** setting on the Preferences page of the administrative console.
Enablement of the notifications allows integration with product tools such as the Rational Application Developer Jython editor to assist you in writing scripts.
The notification type is `websphere.command.assistance.jython.user_name` where `user_name` is the name of the administrative console user.

Results

You have viewed wsadmin scripting commands from the administrative console, optionally logged the commands to a file, and optionally allowed command assistance to emit JMX notifications.

What to do next

You can continue your administration of the administrative console.

Administrative console actions with command assistance:

Using command assistance, you can view wsadmin scripting commands in the Jython language for the last action that runs in the administrative console. This topic lists the administrative console actions that have **wsadmin** commands available in the command assistance option of the Help portlet.

The table lists the components and the actions in the administrative console that have command assistance for a particular component. The administrative console can be an unfederated application server administrative console, a deployment manager administrative console, an administrative agent administrative console, or a job manager administrative console. Listed actions might apply to one, some, or all of the administrative consoles.

Table 4. Console component actions. The following table lists the components and the actions in the administrative console that have command assistance for a particular component.

Component	Action
Applications	<ul style="list-style-type: none"> • List the applications. • Install the application. • Update the complete application. • Start the application. • Stop the application. • Edit the deployment target mapping. • List Structured Query Language in Java (SQLJ) profiles. • Customize and bind SQLJ profiles. • List IBM Optim™ pureQuery Runtime bind files (*). • Bind IBM Optim pureQuery Runtime bind files (*). • Uninstall the application. <p>(*) Command assistance in the administrative console produces wsadmin commands for SQLJ. These commands work with IBM Optim pureQuery Runtime bind files. Command assistance does not produce separate commands for the IBM Optim pureQuery Runtime bind files. However, in the wsadmin environment separate commands exists for you to use when working with IBM Optim pureQuery Runtime bind files.</p>
Web servers	<ul style="list-style-type: none"> • Create a web server. • Delete a web server. • Generate a plug-in configuration. • Propagate the plug-in configuration. • Propagate the key ring for the plug-in configuration. • Start the web server. • Stop the web server. • Terminate the web server.
Channel framework	<ul style="list-style-type: none"> • List the SSL repertoires. • List the Transmission Control Protocol (TCP) endpoints. • List the TCP thread pools. • Delete a chain. • Get the TCP endpoint. • Create a TCP endpoint. • Create a chain.
Node groups	<ul style="list-style-type: none"> • Create a node group. • Remove a node group. • Add a node group member. • Remove a node group member.
Core group bridge	<ul style="list-style-type: none"> • List eligible bridge interfaces. • List tunnel access point groups. • Modify a tunnel access point group. • List tunnel access points. • Modify a tunnel access point. • Modify a peer core group. • List tunnel templates. • Modify a tunnel template.

Table 4. Console component actions (continued). The following table lists the components and the actions in the administrative console that have command assistance for a particular component.

Component	Action
Core group	<ul style="list-style-type: none"> • Create a core group. • Modify a core group. • List the core groups. • Delete a core group. • List the core group servers. • Move a core group server. • Move a core group cluster. • List the core group policies. • Create a core group policy. • Modify a core group policy. • Delete a core group policy. • Create a core group policy match criteria. • Modify a core group policy match criteria. • Delete a core group policy match criteria.
Clusters	<ul style="list-style-type: none"> • List the clusters. • Create a cluster. • Modify the cluster configuration properties. • Modify the cluster runtime properties. • Delete a cluster. • Modify the configuration properties of the domain bootstrap address of a backup cluster. • Modify the runtime properties of the domain bootstrap address of a backup cluster. • List the cluster members. • Create a cluster member. • Modify the cluster member configuration properties. • Delete a cluster member. • List the cluster member templates. • Start a cluster. • Stop a cluster. • Stop a cluster immediately. • Ripple start a cluster. • Start a cluster member. • Restart a cluster member. • Stop a cluster member. • Stop a cluster member immediately. • Terminate a cluster member. • Update the cluster member configuration weight.
Topology	<ul style="list-style-type: none"> • Create an unmanaged node. • Remove an unmanaged node. • Synchronize the nodes. • Stop a node. • Stop a node agent. • Restart a node.

Table 4. Console component actions (continued). The following table lists the components and the actions in the administrative console that have command assistance for a particular component.

Component	Action
Servers	<ul style="list-style-type: none"> • Create an application server. • Delete an application server. • Modify an application server. • List application servers. • Remove application servers. • Start an application server. • Stop an application server. • Create a custom property for an application server. The action is supported for an application server, but not its template. • Modify an existing custom property for an application server. • List custom properties for an application server. The action is supported for an application server, but not its template. • Remove custom properties for an application server. • Create a custom service for an application server. The action is supported for an application server, but not its template. • Modify an existing custom service for an application server. • List custom services for an application server. The action is supported for an application server, but not its template. • Remove custom services for an application server. • Modify application server components. • Modify the Object Request Broker (ORB) service. • Modify the ORB.thread.pool properties for an ORB service. • Modify the thread pool detail for an ORB service.
Servers (continued)	<ul style="list-style-type: none"> • Create a new class loader. The action is supported for an application server, but not its template. • Modify an existing class loader. • List class loaders. The action is supported for an application server, but not its template. • Remove class loaders. • Modify a process definition detail. • Create a new environment entry for an application server. The action is supported for an application server, but not its template. • Modify an existing environment entry for an application server. • List environment entries for an application server. The action is supported for an application server, but not its template. • Remove environment entries for an application server. • Modify Java virtual machine (JVM) configuration properties. • Modify JVM runtime properties.

Table 4. Console component actions (continued). The following table lists the components and the actions in the administrative console that have command assistance for a particular component.

Component	Action
Servers (continued)	<ul style="list-style-type: none"> • Modify process execution properties. • Modify process logs configuration properties. • Modify process logs runtime properties. • Create a new port property for the application server. • Modify an existing port property for the application server. • List ports for the application server. • Remove ports from the application server. • Modify session management properties. • Modify cookie properties. • Modify distributed environment settings. • Modify custom tuning parameters. • Modify custom settings for custom tuning parameters. • Modify database settings. • Create a new thread pool. The action is supported for an application server, but not its template. • Modify an existing thread pool. • List thread pools. The action is supported for an application server, but not its template. • Remove thread pools. • Create an application server template. • Delete an application server template. • List the application server templates.
Servers (continued)	<ul style="list-style-type: none"> • Create a generic server. • Delete a generic server. • Create a new environment entry for a generic server. The action is supported for a generic server, but not its template. • List environment entries for a generic server. The action is supported for a generic server, but not its template. • Modify the Enterprise JavaBeans (EJB) container settings. • Modify the EJB cache settings. • Modify the EJB timer service settings. • Modify the application profiling service. • Modify the internationalization service. • Modify the compensation service. • Modify the object pool service. • Modify the startup beans service. • Modify the ActivitySession service. • Modify the work area service. • Modify the core group service. • Create a work area partition. • Modify a work area partition. • Delete a work area partition.

Table 4. Console component actions (continued). The following table lists the components and the actions in the administrative console that have command assistance for a particular component.

Component	Action
Servers (continued)	<ul style="list-style-type: none"> • Modify web container properties. • Create a custom property. The action is supported for a server, but not its template. • Modify an existing custom property. • List custom properties. The action is supported for a server, but not its template. • Remove custom properties. • Create a new web container transport chain. • Modify an existing web container transport chain. • Remove web container transport chains. • Create an external cache group. • Edit an external cache group. • Edit a denial of service protection. • Modify the default Java persistence application programming interface (API) settings. • View product information. • View installed components. • View installed extensions. • Modify administrative services. • Modify a repository service. • Modify a generic server. • Create a generic server endpoint. • Modify a generic server endpoint. • Remove a generic server endpoint.
Proxy Server	<ul style="list-style-type: none"> • Create a proxy server. • Delete a proxy server. • Modify proxy server security settings. • List proxy server security settings. • List proxy virtual hosts. • Modify a proxy virtual host. • Modify proxy virtual host settings. • List proxy rule expressions. • Modify a proxy rule expression. • List proxy actions. • Modify a caching action. • Modify a compression action. • Modify a header action. • Modify a rewriting action. • Modify a routing action. • Modify a time mapping. • List custom advisors. • Modify a custom advisor. • Modify a custom advisor mapping.

Table 4. Console component actions (continued). The following table lists the components and the actions in the administrative console that have command assistance for a particular component.

Component	Action
Session Initiation Protocol (SIP) container	<ul style="list-style-type: none"> • Modify the Session Initiation Protocol (SIP) container. • Modify the SIP stack. • Modify SIP timers. • Modify the SIP digest authentication. • List SIP application routers. • Modify a SIP application router.
Environment	<ul style="list-style-type: none"> • Create a name space binding for different types. • Edit an EJB name space binding. • Edit a name space binding of indirect lookup. • Edit a name space binding of other context property. • Edit a string name space binding. • Create bootstrap properties. • Edit bootstrap properties. • Create a group of Universal Resource Identifier (URI) patterns. • Edit a group of URI patterns. • Create a shared library. • Edit a shared library. • Modify a shared library. • List shared libraries. • Remove shared libraries. • Create a virtual host. • Edit a virtual host. • Modify a virtual host.
Environment (continued)	<ul style="list-style-type: none"> • List virtual hosts. • Remove virtual hosts. • Create a host alias for a virtual host. • Edit a host alias for a virtual host. • Modify host aliases for a virtual host. • List host aliases for a virtual host. • Remove host aliases for a virtual host. • Create a multi-purpose internet mail extensions (MIME) type. • Edit a MIME type. • Modify a MIME type. • List MIME types. • Remove MIME types. • Create a WebSphere variable. • Edit a WebSphere variable. • Modify a WebSphere variable. • List WebSphere variables. • Remove WebSphere variables.

Table 4. Console component actions (continued). The following table lists the components and the actions in the administrative console that have command assistance for a particular component.

Component	Action
Resources	<ul style="list-style-type: none"> • Create a Java Database Connector (JDBC) provider. • List the JDBC providers. • Modify a JDBC provider. • Delete a JDBC provider. • List the resource adapters. • Install a resource adapter. • Copy a resource adapter. • Upgrade a resource adapter. • Modify a resource adapter. • Delete a resource adapter. • Modify the advanced resource adapter properties of a resource adapter. • Create a custom property. • Modify a custom property. • Delete a custom property. • List the data sources. • Create a data source. • Modify a data source. • Remove a data source. • Modify the connection pool properties of a data source. • Modify the advanced connection pool properties of a data source. • Modify the Websphere Application Server data source properties of a data source.
Resources (continued)	<ul style="list-style-type: none"> • List the WebSphere Application Server Version 4 data sources. • Create a WebSphere Application Server Version 4 data source. • Modify a WebSphere Application Server Version 4 data source. • Delete a Websphere Application Server Version 4 data source. • Modify the connection pool properties of a Websphere Application Server Version 4 data source. • List the Java 2 Connector (J2C) connection factories. • Create a J2C connection factory. • Modify a J2C connection factory. • Delete a J2C connection factory. • Modify the connection pool properties of a J2C connection factory. • Modify the advanced connection pool properties of a J2C connection factory. • Modify the advanced connection factory properties. • List the J2C activation specifications. • Create a J2C activation specification. • Modify a J2C activation specification. • Delete a J2C activation specification. • List the J2C administered objects. • Create a J2C administered object. • Modify a J2C administered object. • Delete a J2C administered object.

Table 4. Console component actions (continued). The following table lists the components and the actions in the administrative console that have command assistance for a particular component.

Component	Action
Resources (continued)	<ul style="list-style-type: none"> • List the schedulers. • Create a scheduler. • Modify a scheduler. • Delete a scheduler. • Create the tables of a scheduler. • Verify the tables of a scheduler. • Drop the tables of a scheduler. • List the object pool managers. • Create an object pool manager. • Modify an object pool manager. • Delete an object pool manager. • Create a custom object pool. • Modify a custom object pool. • Delete a custom object pool. • List the work managers. • Create a work manager. • Modify a work manager. • Delete a work manager. • List the timer managers. • Create a timer manager. • Modify a timer manager. • Delete a timer manager. • Create a mail provider. • Modify a mail provider. • Create a mail session. • Edit a mail session. • Create a protocol provider. • Modify a protocol provider.
Resources (continued)	<ul style="list-style-type: none"> • Create a referenceable. • Modify a referenceable. • Create resource environment entries. • Edit resource environment entries. • Create a resource environment provider. • Edit a resource environment provider. • Create a URL. • Modify a URL. • Create a URL provider. • Modify a URL provider. • Modify an object cache instance. • Modify a servlet cache instance.

Table 4. Console component actions (continued). The following table lists the components and the actions in the administrative console that have command assistance for a particular component.

Component	Action
Transaction services	<ul style="list-style-type: none"> • List the manual transactions. • List the retry transactions. • List the heuristic transactions. • List the imported prepared transactions. • Set the total transaction lifetime timeout. • Set the asynchronous response timeout. • Enable file locking. • Enable transaction coordination authorization. • Set the client inactivity timeout. • Set the maximum transaction timeout.
Security	<ul style="list-style-type: none"> • Enable security. • Validate Lightweight Directory Access Protocol (LDAP) connections. • List SSL configurations. • Get the SSL configuration. • Create an SSL configuration. • Modify the SSL configuration. • Delete an SSL configuration. • List the SSL ciphers. • List the SSL configuration groups. • Create an SSL configuration group. • Delete an SSL configuration group. • Modify an SSL configuration group. • Get the inherited SSL configuration. • List dynamic outbound endpoint SSL configurations. • Create a dynamic outbound endpoint SSL configuration. • Delete a dynamic outbound endpoint SSL configuration. • List the key sets. • Generate a key for a key set. • Create a key set. • Delete a key set. • List the key set groups. • Create a key set group. • Delete a key set group. • Generate keys for the key set group.

Table 4. Console component actions (continued). The following table lists the components and the actions in the administrative console that have command assistance for a particular component.

Component	Action
Security (continued)	<ul style="list-style-type: none"> • List the keystores. • Create a keystore. • Modify a keystore. • Delete a keystore. • Change the keystore password. • Exchange signers. • List the key managers. • Create a key manager. • Delete a key manager. • List the key file aliases. • Create the key reference. • List the trust managers. • Create a trust manager. • Delete a trust manager. • List the certificate authority clients. • Create a certificate authority client. • Modify a certificate authority client. • Delete a certificate authority client. • List the personal certificates. • Get the attributes of a personal certificate. • Get a certificate chain. • Receive a personal certificate. • Create a self-signed certificate. • Create a chained certificate. • Create a certificate authority signed certificate.

Table 4. Console component actions (continued). The following table lists the components and the actions in the administrative console that have command assistance for a particular component.

Component	Action
Security (continued)	<ul style="list-style-type: none"> • Renew a personal certificate. • Revoke a personal certificate. • Replace a personal certificate. • Extract a personal certificate. • Import a personal certificate. • Export a personal certificate. • Delete a personal certificate. • Add a signer certificate. • Extract a signer certificate. • Retrieve signer information from a port. • Retrieve a signer certificate from a port. • Get the properties of a signer certificate. • Delete a signer certificate. • List the signer certificates. • Create a certificate request. • Get a certificate request. • List the certificate requests. • Delete a certificate request. • Extract a certificate request. • Query a certificate request. • List the notifiers. • Create a notifier. • Delete a notifier. • Start the certificate expiration monitor. • Validate the administrative name. • Add a base entry to the realm. • Modify the base entry details. • Configure a new LDAP repository. • Modify an existing LDAP repository configuration. • Delete an existing LDAP repository configuration.

Table 4. Console component actions (continued). The following table lists the components and the actions in the administrative console that have command assistance for a particular component.

Component	Action
Security (continued)	<ul style="list-style-type: none"> • View performance data for the LDAP repository under an LDAP configuration. • Modify the LDAP performance data. • View the LDAP entity types under the LDAP configuration. • Modify the existing LDAP entity types. • View the group attribute definition under the LDAP configuration. • Modify the group attribute definition under the LDAP configuration. • View the member attributes under the LDAP group attribute definitions. • Configure the member attribute details under the LDAP group attribute definitions. • Delete an existing member attribute detail. • View the dynamic member attributes under the LDAP group attribute definitions. • Configure the dynamic member attributes under the LDAP group attribute definitions. • Delete an existing dynamic member attribute detail. • View the list of repositories to manage. • Configure the federated repositories to use a built-in repository. • Remove the built-in repository from the federated repository configuration. • View the federated repository property extension. • Configure the federated repository property extension. • View the federated repository entry mapping repository. • Configure the federated repository entry mapping repository. • View the federated repository supported entity types list. • View the details of a supported entity type. • Modify an existing supported entity type. • View the authentication mechanism and expiration policy for the federated repository user identity. • Get an audit policy. • Modify an audit policy. • List the audit event type filters. • Get an audit event type filter. • Create an audit event type filter. • Modify an audit event type filter.

Table 4. Console component actions (continued). The following table lists the components and the actions in the administrative console that have command assistance for a particular component.

Component	Action
Security (continued)	<ul style="list-style-type: none"> • Delete an audit event type filter. • List the audit service providers. • Get an audit service provider. • Create an audit service provider. • Modify an audit service provider. • Delete an audit service provider. • List the audit factories. • Get an audit factory. • Create an audit factory. • Modify an audit factory. • Delete an audit factory. • List the audit encryption keystores. • Get an audit encryption keystore. • Create an audit encryption keystore. • Modify an audit encryption keystore. • Delete an audit encryption keystore. • Get an audit encryption configuration. • Create an audit encryption configuration. • Modify an audit encryption configuration. • Delete an audit encryption configuration. • Get an audit signing configuration. • Create an audit signing configuration. • Modify an audit signing configuration. • Delete an audit signing configuration. • List the audit notification monitors. • Create an audit notification monitor. • Modify an audit notification monitor, • List the audit notifications. • Create an audit notification • Modify an audit notification. • Delete an audit notification.

Table 4. Console component actions (continued). The following table lists the components and the actions in the administrative console that have command assistance for a particular component.

Component	Action
Security (continued)	<ul style="list-style-type: none"> • List the active security settings. • Set the active administrative security settings. • List the active Rivest Shamir Adleman (RSA) token authorization settings. • Set the active Rivest Shamir Adleman (RSA) token authorization settings. • List all authorization groups. • Create a new authorization group. • Delete an authorization group. • Edit an authorization group. • Get user realm information. • Configure a local OS user realm. • Configure an LDAP user realm. • Configure a custom user realm. • Unconfigure the user realm. • List trusted realms. • Add trusted realms. • Remove trusted realms. • Unconfigure the trusted realm. • Get external authorization provider information. • Set external authorization provider information. • Unconfigure external authorization provider. • List trust association interceptors. • Create a trust association interceptor. • Modify a trust association interceptor. • Delete a trust association interceptor. • Unconfigure the trust association. • Get Common Secure Interoperability (CSI) inbound information. • Set CSI inbound information. • Unconfigure CSI inbound information. • Get CSI outbound information. • Set CSI outbound information. • Unconfigure CSI outbound information.

Table 4. Console component actions (continued). The following table lists the components and the actions in the administrative console that have command assistance for a particular component.

Component	Action
Security (continued)	<ul style="list-style-type: none"> • List Java Authentication and Authorization Service (JAAS) login configurations. • Create JAAS login configurations. • Modify JAAS login configurations. • Delete JAAS login configurations. • Unconfigure JAAS login information. • Configure a JAAS login module. • Delete a JAAS login module. • List JAAS authorization data entries. • Create a JAAS authorization data entry. • Modify the JAAS authorization data entry. • Delete a JAAS authorization data entry. • Get Simple and Protected GSS-API Negotiation Mechanism (SPNEGO) information. • Configure SPNEGO information. • List SPNEGO filters. • Create a SPNEGO filter. • Modify the SPNEGO filter. • Delete a SPNEGO filter. • Create a Kerberos authentication mechanism. • Modify the Kerberos authentication mechanism. • List security domains. • Create a security domain. • Copy the security domain. • Modify the security domain. • Delete the security domain. • Set active security settings. • Unset active security settings. • Set the Lightweight Third-Party Authentication (LTPA) timeout. • Configure programmatic session cookies.

Table 4. Console component actions (continued). The following table lists the components and the actions in the administrative console that have command assistance for a particular component.

Component	Action
Service integration	<ul style="list-style-type: none"> • Create a bus. • Delete a bus. • Add a bus member. • Delete a bus member. • List the bus members. • Delete the messaging engine. • Create a queue. • Create a topic space. • Create an alias destination. • Delete an alias destination. • Create an MQ queue type destination. • Create a foreign destination. • Mediate a destination. • Unmediate a destination. • Delete a destination. • Create a mediation. • Delete a mediation. • Modify a mediation. • Delete a foreign bus. • Create a Java Message Service (JMS) activation specification. • Modify a JMS activation specification. • Create a JMS connection factory. • Modify a JMS connection factory. • Create a JMS queue connection factory. • Modify a JMS queue connection factory. • Create a JMS topic connection factory. • Modify a JMS topic connection factory.

Table 4. Console component actions (continued). The following table lists the components and the actions in the administrative console that have command assistance for a particular component.

Component	Action
Service integration (continued)	<ul style="list-style-type: none"> • Create a JMS queue. • Modify a JMS queue. • Create a JMS topic. • Modify a JMS topic. • Create a JMS provider. • Create a WebSphere MQ server. • Modify a WebSphere MQ server. • Modify a WebSphere MQ server bus member. • Add a permitted transport. • Add a user to a bus connector role. • Add a group to a bus connector role. • Add an inbound port to an inbound service. • Add an outbound port to an outbound service. • Connect an endpoint listener to a service integration bus. • Create an endpoint listener. • Create an inbound service. • Create an outbound service. • Delete an endpoint listener. • Delete an inbound service. • Delete an outbound service. • Disconnect an endpoint listener from a service integration bus. • Publish an inbound service to a Universal Description, Discovery, and Integration (UDDI) registry. • Refresh the Web Services Description Language (WSDL) definition for an inbound service. • Refresh the WSDL definition for an outbound service. • Remove an inbound port. • Remove an outbound port. • Set the default outbound port for an outbound service. • Remove an inbound service from a UDDI registry. • Add a target service for routing from a gateway service. • Create a gateway service. • Create a proxy service deployed to a gateway instance. • Delete a gateway service. • Delete a gateway instance. • Delete a proxy service. • Remove a target service. • Create a WS-Notification (WSN) service.

Table 4. Console component actions (continued). The following table lists the components and the actions in the administrative console that have command assistance for a particular component.

Component	Action
Service integration (continued)	<ul style="list-style-type: none"> • Modify a bus. • Modify a foreign bus connection. • List foreign bus connections. • List bus messaging engines. • List messaging engines on a particular server. • List bus mediations. • Modify a messaging engine. • Create a service integration bus (SIB) link on a messaging engine. • Modify a foreign bus connection (an indirect link). • Modify a foreign bus connection (a SIB link). • List bus destinations. • List queue points (point-to-point messaging) for a destination. • List queue mediation points for a destination. • List topic mediation points for a destination. • Create a new context property. • Edit a bus. • Create a replication domain. • Edit a replication domain.
System administration	<ul style="list-style-type: none"> • Edit the file synchronization service. • Edit the file transfer service. • Modify a cell. • Modify the deployment manager. • Modify a node. • Modify a node agent. • Modify a Java Management Extensions (JMX) connector. • Modify extension MBean providers. • Modify node groups.
Web services	<ul style="list-style-type: none"> • List the service clients in a cell. • List the service clients in an application. • List the service providers in a cell. • List the service providers in an application. • Start the service provider listener. • Stop the service provider listener.

Table 4. Console component actions (continued). The following table lists the components and the actions in the administrative console that have command assistance for a particular component.

Component	Action
Web services policy sets	<ul style="list-style-type: none"> • List the policy sets in the system. • Get the attributes for a policy set. • Set the attributes of a specified policy set. • Create a new policy set. • Copy a policy set to create a new policy set. • Delete a policy set. • Update the attributes of a policy set. • Export a policy set from an archive for use in a client environment or a server. • Get the policy set attachments for a given resource. • List the policy set attachments for services providers. • List the policy set attachments for services clients in an application. • Create a new policy set attachment for a resource. • Delete a policy set attachment from a resource. • List the applications to which a given policy set is attached. • Delete all attachments for a policy set. • Transfer all attachments from one policy set to another.
Web services policies	<p>Attention: Before you use the generated wsadmin command for policies, see the documentation for the PolicySetManagement command group for the AdminTask object.</p> <ul style="list-style-type: none"> • Create a policy. • Add a policy to a policy set. • Delete a policy from a policy set. • List the names of existing policies. • Get the attributes for a policy. • Update the configuration of a policy. • Get the value for a named policy attribute. • Set the value for a named policy attribute.
Web services bindings	<ul style="list-style-type: none"> • Get the binding configuration for a specified policy for a policy set attachment. • Set the binding for a policy set attachment. • Set and update the binding configuration for a specified policy for a policy set attachment.
Web services trust service	<ul style="list-style-type: none"> • List the local names of all the configured token providers. • Query the trust service for the local name of the default token provider. • Update configuration data for a token provider. • Delete custom properties from a token provider configuration. • Assign a token that is issued when requesting access to a specific end point. • List the assigned endpoints for a token provider. • Query the trust service for the token provider assigned to a specified endpoint. • Unassign an endpoint from its token provider. • Refresh trust service. • Query the trust service for a list of assigned endpoints.

Table 4. Console component actions (continued). The following table lists the components and the actions in the administrative console that have command assistance for a particular component.

Component	Action
Web Services Security distributed cache	<ul style="list-style-type: none"> • Get the Web Services Security distributed cache configuration. • Get the Web Services Security distributed cache configuration custom properties. • Get the defined cell level data sources. • Set the Web Services Security distributed cache configuration. • Set the Web Services Security distributed cache configuration custom properties.
Business-level applications	<ul style="list-style-type: none"> • List the assets. • Add an asset to the repository. • Export an asset. • Delete an asset . • List the business-level applications. • Create a new business-level application. • Add an asset to a business-level application. • Delete a business-level application. • Edit a business-level application. • Start a business-level application. • Stop a business-level application. • Edit a composition unit.
Performance Monitoring Infrastructure (PMI) and Request metrics	<ul style="list-style-type: none"> • List the PMI configuration. • Modify the PMI configuration. • Modify the PMI parameters at run time. • Edit the request metrics filter. • Edit the request metrics. • Create a request metrics filter value. • Edit a request metrics filter value. • Delete the request metrics filter.
Performance advisors	<ul style="list-style-type: none"> • Re-initialize the Runtime Performance Advisor tool. • Set the Runtime Performance Advisor tool attributes.
Portlets and portlet containers	<ul style="list-style-type: none"> • View the portlet deployment descriptor. • Modify portlet container settings. • Create a custom property. • Modify an existing custom property. • List custom properties. • Remove custom properties. • Enable PMI for portlets. • Enable request metrics for portlets.

Table 4. Console component actions (continued). The following table lists the components and the actions in the administrative console that have command assistance for a particular component.

Component	Action
DataPower®	<ul style="list-style-type: none"> • Add an appliance to the DataPower appliance manager. • Add a firmware version to the DataPower appliance manager. • Add a managed set to the DataPower appliance manager. • Copy a DataPower appliance manager managed domain version to a new managed set. • Copy a DataPower appliance manager managed settings version to a new managed set. • Export the configuration and versions of a DataPower appliance manager. • Get the IDs of all the DataPower appliance manager appliances. • Get the names of all the domains on a DataPower appliance. • Get the IDs of all the DataPower appliance manager firmware. • Get the IDs of all the DataPower appliance manager firmware versions. • Get the IDs of all the appliances in a DataPower appliance manager managed set. • Get the IDs of all the domains in a DataPower appliance manager managed set. • Get the IDs of all the domain versions in a DataPower appliance manager managed set. • Get the IDs of the managed sets using a firmware version. • Get the IDs of all the settings versions in a DataPower appliance manager managed set. • Get the IDs of all the DataPower appliance manager managed sets. • Get the IDs of all the DataPower appliance manager tasks. • Get the firmware that best matches the parameters. • Get a DataPower appliance manager appliance. • Get a DataPower appliance manager firmware. • Get a DataPower appliance manager firmware version. • Get a DataPower appliance manager managed domain.

Table 4. Console component actions (continued). The following table lists the components and the actions in the administrative console that have command assistance for a particular component.

Component	Action
DataPower (continued)	<ul style="list-style-type: none"> • Get a DataPower appliance manager managed domain version. • Get a DataPower appliance manager managed settings. • Get a DataPower appliance manager managed settings version. • Get a DataPower appliance manager managed set. • Get the properties of the DataPower appliance manager. • Get a DataPower appliance manager task. • Import the configuration and versions of the DataPower appliance manager. • Tell the DataPower appliance manager to add the appliance to a managed set and start managing the appliance. • Tell the DataPower appliance manager to add the domain to a managed set and start managing the appliance. • Purge a DataPower appliance manager task. • Remove an appliance from the DataPower appliance manager. • Remove a firmware version from the DataPower appliance manager. • Remove a managed domain version from the DataPower appliance manager. • Remove a managed settings version from the DataPower appliance manager. • Remove a managed set from the DataPower appliance manager. • Modify a DataPower appliance manager appliance. • Modify a DataPower appliance manager firmware version. • Modify a DataPower appliance manager managed domain. • Modify a DataPower appliance manager managed domain version. • Modify a DataPower appliance manager managed settings. • Modify a DataPower appliance manager managed settings version. • Modify a DataPower appliance manager managed set. • Modify the DataPower appliance manager.
DataPower (continued)	<ul style="list-style-type: none"> • Manually synchronize a DataPower appliance manager managed set. • Tell the DataPower appliance manager to remove the appliance from its managed set and stop managing the appliance. • Tell the DataPower appliance manager to remove the domain from its managed set and stop managing the appliance. • Get the status of the DataPower appliance manager. • Shut down the DataPower appliance manager.
Replication Domains	<ul style="list-style-type: none"> • List data replication domain members.

Table 4. Console component actions (continued). The following table lists the components and the actions in the administrative console that have command assistance for a particular component.

Component	Action
Job manager	<ul style="list-style-type: none"> • Query targets. • Get target properties. • Query target resources. • Get target resource properties. • List groups of targets. • Create a group of targets. • Delete a group of targets. • Get information about groups of targets. • Modify information about groups of targets. • List members of a group of targets. • Add a member to a group of targets. • Delete a member from a group of targets. • Query jobs. • Get job status. • Get job target status. • Get job target history. • Delete a job. • Suspend a job. • Resume a job.
Users and groups	<ul style="list-style-type: none"> • Add a new user. • Modify an existing user. • List users. • Remove users. • Add a new group. • Modify an existing group. • List groups. • Remove groups. • List registry users. • List registry groups.

Changing the console session expiration

Run this JACL script to set how long Integrated Solutions Console can be used until the login session expires.

About this task

The following JACL script serves as an example of how to set the duration that an Integrated Solutions Console can be used until the login session expires. Other scripting types, such as JYTHON, could be used.

Procedure

1. Copy the following script into a file.

```
set dep [AdminConfig getid /Deployment:isclite/]
set appDep [AdminConfig list ApplicationDeployment $dep]
set sesMgmt [AdminConfig list SessionManager $appDep]
```

```

# check if existing sesMgmt there or not, if not then create a new one, if exist then modify it
if {$sesMgmt == ""} {
  # get applicationConfig to create new SessionManager
  set appConfig [$AdminConfig list ApplicationConfig $appDep]
  if {$appConfig == ""} {
    # create a new one
    set appConfig [$AdminConfig create ApplicationConfig $appDep {}]
    # then create a new SessionManager using new Application Config just created
    set sesMgmt [$AdminConfig create SessionManager $appConfig {}]
  } else {
    # create new SessionManager using the existing ApplicationConfig
    set sesMgmt [$AdminConfig create SessionManager $appConfig {}]
  }
}

# get tuningParams config id
set tuningParams [$AdminConfig showAttribute $sesMgmt tuningParams]
if {$tuningParams == ""} {
  # create a new tuningParams
  $AdminConfig create TuningParams $sesMgmt {{invalidationTimeout <timeout value>}}
} else {
  #modify the existing one
  $AdminConfig modify $tuningParams {{invalidationTimeout <timeout value>}}
}

# saving the configuration changes
$AdminConfig save

```

2. Change the *<timeout value>* on the two lines of this sample to the new session expiration value. This number specifies the number of minutes the console preserves the session during inactivity.
3. Save the file to any directory using, for example, the filename `timeout.jacl`.
4. Start the `wsadmin` scripting client from the `<WAS-install>/profiles/<profile_name>/bin` directory.
5. Issue the following command.

```
wsadmin -f <path to jacl file>/timeout.jacl
```

Changing the class loader order of the console module deployed in Integrated Solutions Console

Run this JACL script to change the class loader order of the console module deployed in the Integrated Solutions Console.

About this task

The following JACL script serves as an example of how to change the class loader order of the console module deployed in the Integrated Solutions Console. Other scripting types, such as JYTHON, could be used.

Procedure

1. Copy the following script into a file.

```

set app [$AdminConfig getid /Deployment:isclite/]
set webModules [$AdminConfig list WebModuleDeployment $app]

foreach webModule $webModules {
  set uri [$AdminConfig showAttribute $webModule uri]
  if {$uri == "<WAR_NAME>"} {
    #modify the classloader for <WAR_NAME>
    set cl [$AdminConfig list Classloader $webModule]
    # check if the classloader exist
    if {$cl == ""} {

```

```

        # create a new one with the appropriate mode
        $AdminConfig create Classloader $webModule {{mode <MODE>}}
    } else {
        # modify the existing one
        $AdminConfig modify $cl {{mode <MODE>}}
    }
}

# save the configuration change
$AdminConfig save

```

2. Change the <WAR_NAME> on the two lines of this sample to the name of the console module file deployed in the Integrated Solutions Console which class loader order you want to change.
3. Change the <MODE> on the two lines of this sample to PARENT_LAST or PARENT_FIRST as required.
4. Save the file to any directory using, for example, the file name classloaderorder.jacl.
5. Start the wsadmin scripting client from the <WAS-install>/profiles/<profile_name>/bin directory.
6. Issue the following command.

```
wsadmin -f <path to jacl file>/classloader.jacl
```

Getting started with wsadmin scripting

Scripting is a non-graphical alternative that you can use to configure and manage WebSphere Application Server.

Before you begin

Verify that user IDs that run WebSphere Application Server for z/OS® scripts, including server, administrator and client user IDs, run with the LANG and LC_ALL environment variables set to the same locale based on code page IBM-1047. Settings based on any other code page might cause the scripts to fail. See the "Changing the Locale in the Shell" topic in *UNIX System Services User's Guide* for more information.

About this task

The WebSphere Application Server wsadmin tool provides the ability to run scripts. The wsadmin tool supports a full range of product administrative activities.

The following figure illustrates the major components involved in a wsadmin scripting solution:

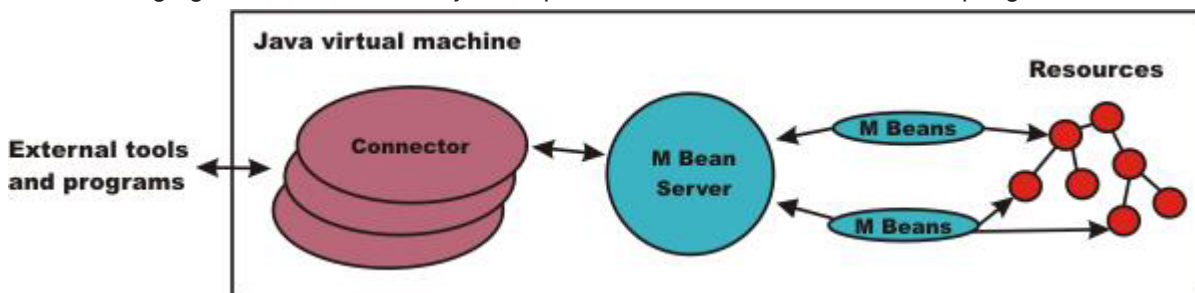


Figure 1: A WebSphere Application Server scripting solution

The wsadmin tool supports two scripting languages: Jacl and Jython. Five objects are available when you use scripts:

- **AdminControl:** Use to run operational commands.

- **AdminConfig:** Use to run configurational commands to create or modify WebSphere Application Server configurational elements.
- **AdminApp:** Use to administer applications.
- **AdminTask:** Use to run administrative commands.
- **Help:** Use to obtain general help.

The scripts use these objects to communicate with MBeans that run in WebSphere Application Server processes. MBeans are Java objects that represent Java Management Extensions (JMX) resources. JMX is an optional package addition to Java 2 Platform Standard Edition (J2SE). JMX is a technology that provides a simple and standard way to manage Java objects.

Important: Some wsadmin scripts, including the AdminApp install, AdminApp update, and some AdminTask commands, require that the user ID under which the server is running must have read permission to the files that are created by the user that is running wsadmin scripting. For example, if the application server is running under user1, but you are running wsadmin scripting under user2, you might encounter exceptions involving a temporary directory. When user2 runs wsadmin scripting to deploy an application, a temporary directory for the enterprise application archive (EAR) file is created. However, when the application server attempts to read and unzip the EAR file as user1, the process fails. It is not recommended that you set the umask value of the user that is running wsadmin scripting to 022 or 023 to work around this issue. This approach makes all of the files that are created by the user readable by other users. To resolve this issue, consider the following approaches based on your administrative policies:

- Run wsadmin scripting with the same user ID as the user that runs the deployment manager or application server. A root user can switch the user ID to complete these actions.
- Set the group ID of the user that is running the deployment manager or application server to be the same group ID as the user that is running wsadmin scripting. Also, set the umask value of the user that is running the wsadmin scripting to be at least a umask 027 value so that files that are created by the wsadmin scripting can be read by members of the group.
- Run wsadmin scripting from a different machine. This approach forces files to be transferred and bypasses the file copy permission issue.

To perform a task using scripting, you must first perform the following steps:

Procedure

1. Choose a scripting language. The wsadmin tool only supports Jacl and Jython scripting languages. Jacl is the language specified by default. If you want to use the Jython scripting language, use the `-lang` option or specify it in the `wsadmin.properties` file.
2. Start the wsadmin scripting client interactively, as an individual command, in a script, or in a profile.

What to do next

Before you perform any task using scripting, make sure that you are familiar with the following concepts:

- Java Management Extensions (JMX)
- WebSphere Application Server configuration model
- wsadmin tool
- Jacl syntax or Jython syntax
- Scripting objects

Optionally, you can customize your scripting environment. For more information, see Administrative properties for using wsadmin scripting.

After you become familiar with the scripting concepts, choose a scripting language, and start the scripting client, you are ready to perform tasks using scripting.

What is new for scripted administration (wsadmin)

This topic highlights what is new or changed for users who are going to customize, administer, monitor, and tune production server environments using the wsadmin tool.

The Deprecated, stabilized, and removed features topic describes features that are being replaced or removed in this or future releases.

Improved administrative scripting features

Enhancements to AdminTask command support for managing configurations using properties files

Version 7 introduced system configuration using properties files. Using commands in the PropertiesBasedConfiguration group, you can copy configuration properties from one environment to another, troubleshoot configuration issues, and apply one set of configuration properties across multiple profiles, nodes, cells, servers, or applications.

Enhancements to configuration using properties files for Version 8.0 include the following:

- Support for web services endpoint URL fragment properties files

For more information, see Using properties files to manage system configuration and PropertiesBasedConfiguration command group for the AdminTask object using wsadmin scripting.

Use the commands and parameters in the AdminSDKCmnds group for the AdminTask object to perform the following actions:

- List software development kits that are not used by a node.
- Get or set the default SDK for a node.
- Get or set an SDK for a server.

For more information, see AdminSDKCmnds command group for the AdminTask object.

AdminSDKCmnds command group for the AdminTask object

getAvailableSDKsOnNode and getSDKPropertiesOnNode commands in the ManagedObjectMetadata command group for the AdminTask object

Use the getAvailableSDKsOnNode and getSDKPropertiesOnNode commands and parameters in the ManagedObjectMetadata group for the AdminTask object to perform the following actions:

- List software development kits that are installed with the product and available for use by a node.
- List software development kit (SDK) properties.

For more information, see ManagedObjectMetadata command group for the AdminTask object.

listServices command for information about web service references

Use the `listServices` command to learn about service references. You can use the `serviceRef` property with the `queryProps` parameter with the `listServices` command to query all service references or a specific service reference. This parameter is only applicable for service clients. If you specify an asterisk (*) as a wildcard as the name of the service reference, all of the service references for the matching service client are returned. You can also query a specific service reference name by specifying the name of the service reference that you want. To return detailed service reference information for endpoints and operations, specify the `expandResource` property.

For more information, see [Querying web services using wsadmin scripting](#).

Security cookies set as HTTPOnly resist cross-site scripting attacks

Use the `HttpOnly` browser attribute to prevent client side applications (such as Java scripts) from accessing cookies to prevent some cross-site scripting vulnerabilities. The attribute specifies that LTPA and WASReqURL cookies include the `HTTPOnly` field.

For more information, see [Single sign-on settings](#).

runCommand, manageprofiles, installIM, updateIM, manageOfferings, and findIMDataLocation parameters of the submitJob command for the AdminTask object

Use new parameters of the `submitJob` command for the `AdminTask` object to use a job manager to perform the following actions:

- Run a command on a remote host.
- Create, augment, or delete a profile.
- Install Installation Manager instances, update Installation Manager with a repository, manage Installation Manager offerings, and install WebSphere Application Server products.

For more information, see [Administrative job types using wsadmin scripting](#).

Overview and new features for scripting the application serving environment

Use the links provided in this topic to learn about the administrative features.

“What is new for scripted administration (wsadmin)” on page 58

This topic provides an overview of new and changed features for administrative scripting and the `wsadmin` tool.

Introduction: Administrative scripting (wsadmin)

This topic provides an introduction to administrative scripting and the `wsadmin` tool.

Using administrative programs (JMX)

This topic describes how to use Java application programming interfaces (APIs) to administer WebSphere Application Server and to manage your applications.

Before you begin

You can administer WebSphere Application Server and your applications through tools that come with the product or through programming with the Java APIs.

The wsadmin scripting tool, the administrative console, and the administrative command-line tools come with the product. These administrative tools provide most of the functions that you need to manage the product and the applications that run in WebSphere Application Server. You can use the command-line tools from automation scripts to control the servers. Scripts that are written for the wsadmin scripting tool offer a wide range of possible custom solutions that you can develop quickly.

Investigate these tools with the Java APIs to determine the best ways to administer WebSphere Application Server and your applications. For information on the Java APIs, view the application programming interfaces documentation.

- **No action required for WAS JMX APIs:** Each Java virtual machine (JVM) in WebSphere Application Server includes an embedded implementation of Java Management Extensions (JMX). In Application Server, Version 5, the JVMs contain an implementation of the JMX 1.0 specification. In Application Server, Version 6.0 and later, the JVMs contain an implementation of the JMX 1.2 specification. The JMX 1.0 implementation used in Version 5 is the TMX4J package that IBM Tivoli® products supply. The JMX 1.2 specification used in Version 6.0 and later is the open source mx4j package. The JMX implementation change across the releases does not affect the behavior of the JMX MBeans in the Application Server. No Application Server administrative application programming interfaces (APIs) are altered due to the change from the JMX V1.0 specification to the JMX V1.2 specification.
- **Action might be required for custom MBeans:** The JMX V1.2 specification is compatible with the earlier JMX V1.0 specification. However, you might need to migrate custom MBeans that are supplied by products other than the Application Server from Version 5 to Version 6.0 and later. The primary concern for these custom MBeans is related to the values that are used in key properties of the JMX ObjectName class for the MBean. The open source mx4j implementation more stringently enforces property validation according to the JMX 1.2 specification. Test the custom MBeans that you deployed in Version 5 in Version 6.0 and later, to ensure compatibility. Full details of the JMX V1.2 specification changes from the JMX V1.0 specification are available in the JMX 1.2 specification.

About this task

WebSphere Application Server supports access to the administrative functions through a set of Java classes and methods. You can write a Java program that performs any of the administrative features of the WebSphere Application Server administrative tools. You can also extend the basic WebSphere Application Server administrative system to include your own managed resources.

You can prepare, install, uninstall, edit, and update applications through programming. Preparing an application for installation involves collecting various types of WebSphere Application Server-specific binding information to resolve references that are defined in the application deployment descriptors. This information can also be modified after installation by editing a deployed application. Updating consists of adding, removing or replacing a single file or a single module in an installed application, or supplying a partial application that manipulates an arbitrary set of files and modules in the deployed application. Updating the entire application uninstalls the old application and installs the new one. Uninstalling an application removes it entirely from the WebSphere Application Server configuration.

Perform any or all of the following tasks to manage WebSphere Application Server and your Java Platform, Enterprise Edition (Java EE) applications through programming.

Procedure

- Create a JMX remote client program by using the JMX remote API (JSR 160)..
This topic describes how to develop a JMX remote program that uses the JMX remote API (JSR 160) to access the WebSphere Application Server administrative system.
- Create a custom Java administrative client program using the Java administrative APIs.
This topic describes how to develop a Java program that uses the WebSphere Application Server administrative APIs to access the administrative system of WebSphere Application Server.
- Extend the WebSphere Application Server administrative system with custom MBeans.

This topic describes how to extend the WebSphere Application Server administration system by supplying and registering new JMX MBeans in one of the Application Server processes. In this case, you can use the administrative classes and methods to add newly managed objects to the administrative system.

- Deploy and manage a custom Java administrative client program for use with multiple Java Platform, Enterprise Edition application servers.

This topic describes how to connect to a Java EE server, and how to manage multiple vendor servers.

Results

Depending on which tasks you complete, you have created your own administrative program, extended the WebSphere Application Server administrative console, connected and managed vendor servers, or managed your applications through programming.

What to do next

You can continue to administer WebSphere Application Server and your applications through programming or in combination with the tools that come with the WebSphere Application Server.

Java Management Extensions (JMX) for WebSphere Application Server

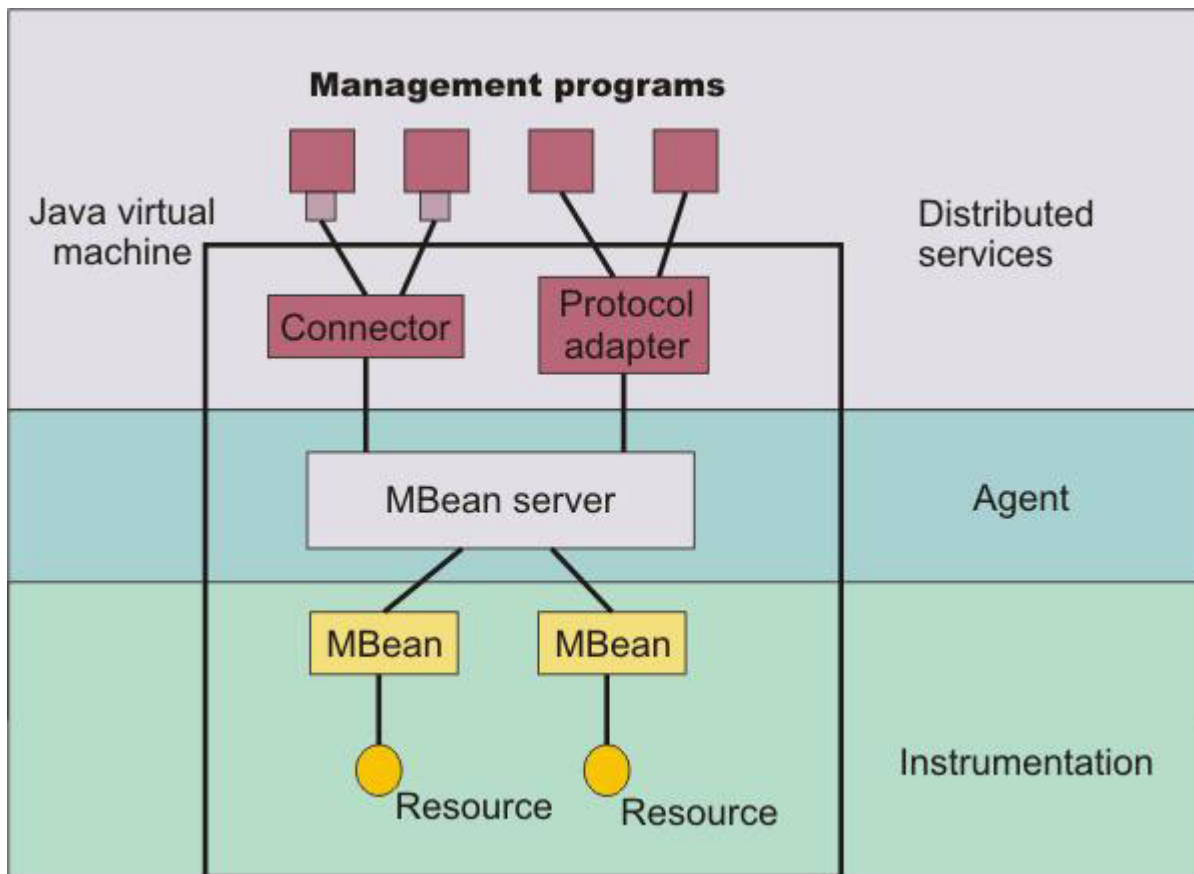
This topic gives an overview of Java Management Extensions (JMX) in general and how this standard applies to WebSphere Application Server.

Java Management Extensions overview

Java Management Extensions (JMX) is the Java standard for managing application resources. The management architecture that is defined by JMX is divided into three levels:

- The bottom level is management instrumentation. Each manageable resource is described by an interface that specifies the attributes it has, the operations it supports, and the notifications it sends. This resource is a managed bean (MBean).
- The middle level is the management agent. Each managed process contains a JMX agent that includes an MBean server, which provides a registry and access point for MBeans. Management clients must use the MBean server to access the registered MBeans.
- The top level of the architecture is defined by JMX Remote application programming interface (API) (JSR 160). JSR 160 uses Remote Method Invocation over Internet Inter-ORB Protocol (RMI-IIOP), but is not interoperable with the RMI connector. The RMI, SOAP/HTTP, and SOAP/HTTPS connectors were created before the JSR160 specification and are supported. The Inter-Process Communications (IPC) connector is also supported.

The top level of the architecture is the distributed services level, and its role is to facilitate remote access to JMX agents. This task is accomplished through connectors, which provide a protocol-independent, location-transparent, client-side interface to the MBean server (for example, a Remote Method Invocation (RMI) connector), or protocol adapters, which provide protocol-specific, server-side access to the MBean server (for example, an HTTP adapter).




Java Management Extensions in WebSphere Application Server

Java Management Extensions (JMX) is at the core of Application Server administration capabilities. The application server contains a JMX agent. All of the system components are defined as MBeans. The JMX agent in Application Server supports the following connectors: JSR160RMI, Remote Method Invocation/Internet Inter-ORB Protocol (RMI/IIOP), Simple Object Access Protocol/Hypertext Transfer Protocol (SOAP/HTTP), Simple Object Access Protocol/Hypertext Transfer Protocol Secure (SOAP/HTTPS), and Inter-Process Communications (IPC), which provides remote access to the server resources. All of the administration tools included with Application Server use these JMX facilities to accomplish their functions.

In a stand-alone Application Server installation, servers exist and are administered individually. An administrative client connects directly to the Application Server in this environment.

In a WebSphere Application Server, Network Deployment installation, a hierarchical topology groups application servers within nodes and groups nodes within a cell. Administrative servers exist at the node level (node agents) and at the cell level (the deployment manager), and act as aggregation points for the administrative services in the subordinate servers.

MBeans in all servers on a node are visible through that node agent, and MBeans in all nodes are visible through the deployment manager. Therefore, by connecting to the deployment manager, you can invoke operations, can get and set attributes, and can receive notifications for any MBean in the cell. Application Server provides an AdminService class that reflects the standard JMX MBeanServer interface, and wraps the MBeanServer interface so that it takes part in implementing this distributed management functionality.



Creating a custom Java administrative client program using WebSphere Application Server administrative Java APIs

This section describes how to develop a Java program for accessing the WebSphere Application Server administrative system by using the product administrative application programming interfaces (APIs).

Before you begin

This task assumes a basic familiarity with Java Management Extensions (JMX) API programming. For information on the Java APIs, view the application programming interfaces documentation.

About this task

When you develop and run administrative clients that use various JMX connectors and that have security enabled, use the following guidelines. When you follow these guidelines, you guarantee the behavior among different implementations of JMX connectors. Any programming model that strays from these guidelines is unsupported.

1. Create and use a single administrative client before you create and use another administrative client.
2. Create and use an administrative client on the same thread.
3. Use one of the following ways to specify a user ID and password to create a new administrative client:
 - Specify a default user ID and password in the property file.
 - Specify a user ID and password other than the default. Once you create an administrative client with a nondefault user ID and password, specify the nondefault user ID and password when you create subsequent administrative clients.

Procedure

1. Develop an administrative client program.
2. If your administrative client uses Simple Object Access Protocol (SOAP) as its Connector, you must configure Java Secure Socket Extension (JSSE) as your transport layer. The product uses supports Secure Sockets Layer (SSL) and Transport Layer Security (TLS) with the Java Secure Sockets Extension (JSSE) and System SSL packages. Create digital certificates for the user ID used to run your administrative client. If you want to use System Authorization Facility (SAF) to create digital certificates and store them in a SAF keyring, refer to Defining SSL Security for Client Outbound Requests. (You can save the name of the keyring you create for use in the next step.)
3. Update the `soap.client.props` file in the `profile_root/properties` directory being used by your administrative client with the name of the SAF keyring. Refer to the second step in Using System Authorization Facility keyrings with Java Secure Sockets Extension for directions on updating the `soap.client.props` file.
4. Build the administrative client program.
Compile it with the `javac` command and provide the location of the necessary JAR files in the classpath argument.

For example, if your installation directory is `/DeploymentManager` a typical command would look like the following example:

```
javac -extdirs "$JAVA_HOME/lib/ext;  
/DeploymentManager/classes;/DeploymentManager/lib;  
/DeploymentManager/lib/ext" MyAdminClient.java
```

(The previous command is split on multiple lines for publication.)

5. Run the administrative client program.
Run the administrative client program by setting up the run-time environment so that the program can find all of the prerequisites. Many of the batch or script files in the `bin` directory under the installation root perform a similar function. The following is an example of a batch file that runs an administrative client program named `MyAdminClient` follows:

```
@echo off  
  
binDir=`dirname "$@"`  
. "$binDir/setupCmdLine.sh"  
  
"$JAVA_HOME/bin/java" "$CLIENTSOAP" "-Dwas.install.root=$WAS_HOME"  
"-Dwas.repository.root=$CONFIG_ROOT"  
-Dcom.ibm.CORBA.BootstrapHost=$COMPUTERNAME  
"-Djava.ext.dirs=$JAVA_HOME/lib/ext;$WAS_HOME/classes;  
$WAS_HOME/lib;$WAS_HOME/lib/ext" MyAdminClient $@
```

(The contents of the previous batch file is split on multiple lines for publication.)

Developing an administrative client program

This topic describes how to develop an administrative client program that utilizes WebSphere Application Server administrative application programming interfaces (APIs) and Java Management Extensions (JMX).

About this task

Product administrative APIs provide control of the operational aspects of your distributed system as well as the ability to update your configuration. For information about the `AdminClient` interface, view the application programming interfaces documentation.

This topic also demonstrates examples of MBean operations. For information on MBean programming, see MBean Java API documentation.

Procedure

1. Create an `AdminClient` instance.

An administrative client program needs to invoke methods on the AdminService object that is running in the deployment manager or the application server in the base installation. The AdminClient class provides a proxy to the remote AdminService object through one of the supported Java Management Extensions (JMX) connectors.

- The following example shows how to create an AdminClient instance for the Simple Object Access Protocol (SOAP) connector:

```
Properties connectProps = new Properties();
connectProps.setProperty(
AdminClient.CONNECTOR_TYPE, AdminClient.CONNECTOR_TYPE_SOAP);

connectProps.setProperty(AdminClient.CONNECTOR_HOST, "localhost");
connectProps.setProperty(AdminClient.CONNECTOR_PORT, "8879");
connectProps.setProperty(AdminClient.USERNAME, "test2");
connectProps.setProperty(AdminClient.PASSWORD, "user24test");
AdminClient adminClient = null;
try
{
    adminClient = AdminClientFactory.createAdminClient(connectProps);
}
catch (ConnectorException e)
{
    System.out.println("Exception creating admin client: " + e);
}
}
```

- a. Set up a Properties object.

The example sets up a Properties object with the properties that are required to get to your server. In this case, you use the SOAP connector to reach the server; for the connector type, use the value: AdminClient.CONNECTOR_TYPE_SOAP.

- b. Set the port number on which the server SOAP connector is listening.

In a single server installation, the default port number for the application server SOAP connector is 8880. In a WebSphere Application Server, Network Deployment installation, the default port number for the deployment manager SOAP connector is 8879.

- c. After the connection properties are set, use the AdminClientFactory class and the Properties object to create an AdminClient object that is connected to your chosen server.

Depending on factors such as your desired protocol and security environment, you might need to set other properties. For example, if you enable security for your application client program, include the javax.net.ssl.* properties. For more detailed information about the AdminClient interface, the javax.net.ssl.* properties, and additional creation examples, refer to the AdminClient interface in the application programming interfaces documentation.

- The following example shows how to create an AdminClient instance for the Remote Method Invocation (RMI) connector. Some commands are split on multiple lines for printing purposes.

```
Properties connectProps = new Properties();
connectProps.setProperty(AdminClient.CONNECTOR_TYPE, AdminClient.CONNECTOR_TYPE_RMI);
connectProps.setProperty(AdminClient.CONNECTOR_HOST, "localhost");
connectProps.setProperty(AdminClient.CONNECTOR_PORT, "2809");
connectProps.setProperty(AdminClient.USERNAME, "test2");
connectProps.setProperty(AdminClient.PASSWORD, "user24test");
System.setProperty("com.ibm.CORBA.ConfigURL",
"file:C:/AA/cf010839.26/profiles/AppSrv02/properties/sas.client.props");
System.setProperty("com.ibm.SSL.ConfigURL",
"file:C:/AA/cf010839.26/profiles/AppSrv02/properties/ssl.client.props");
AdminClient adminClient = null;
try
{
    adminClient = AdminClientFactory.createAdminClient(connectProps);
}
catch (ConnectorException e)
{
    System.out.println("Exception creating admin client: " + e);
}
}
```

Note: When you use the `createAdminClient` method within application code that runs on an application server, such as within servlets and JavaServer Pages (JSP) files, you must set the `CACHE_DISABLED` property to true. For example:

```
connectProps.setProperty(AdminClient.CACHE_DISABLED, "true");
```

- a. Set up a Properties object.

The example sets up a Properties object with the properties that are required to get to your server. In this case, you use the Remote Method Invocation connector to reach the server; for the connector type, use the value: `AdminClient.CONNECTOR_TYPE_RMI`.

- b. Set the port number on which the server RMI connector is listening.

In a single server installation, the default port number for the application server RMI connector is 2809. In a WebSphere Application Server, Network Deployment installation, the default port number for the deployment manager RMI connector is 9809.

- c. After the connection properties are set, use the `AdminClientFactory` class and the Properties object to create an `AdminClient` object that is connected to your chosen server.

Depending on factors such as your desired protocol and security environment, you might need to set other properties. For example, if you enable security for your application client program, you need to set a system property to point to the `ssl.client.props` file and the `sas.client.props` file. If you run on a local machine you can point to the actual location. If you run on a remote machine, you can copy these properties files from the server machine and put them anywhere you want, specifying the path to where you put the files.

You can specify a user name and password inside the `sas.client.props` file. When you do, specify `com.ibm.CORBA.loginSource=properties`. If you want to set the user name and password inside your client program, specify `com.ibm.CORBA.loginSource=none` in the `sas.client.props` file.

2. Find an MBean.

When you obtain an `AdminClient` instance, you can use it to access managed resources in the administration servers and application servers. Each managed resource registers an MBean with the `AdminService` through which you can access the resource. The MBean is represented by an `ObjectName` instance that identifies the MBean. An `ObjectName` instance consists of a domain name followed by an unordered set of one or more key properties. The syntax for the domain name follows:

```
[domainName]:property=value[,property=value]*
```

For WebSphere Application Server, the domain name is `WebSphere` and the key properties defined for administration are as follows:

Table 5. Key property descriptions. Key properties include types, name, cell, node, and process.

type	The type of MBean. For example: <code>Server</code> , <code>TraceService</code> , Java virtual machine (JVM).
name	The name identifier for the individual instance of the MBean.
cell	The name of the cell that the MBean is running.
node	The name of the node that the MBean is running.
process	The name of the process that the MBean is running.

Some MBeans in WebSphere Application Server use additional key properties. An MBean without key properties can be registered with the MBean server in a WebSphere Application Server process. However, such an MBean cannot participate in the distributed enhancements that the product adds, for example, request routing, distributed event notification, and so on.

If you know the complete set of key properties for an `ObjectName` instance, you can use it to find the MBean it identifies. However, finding MBeans without having to know all of their key properties is usually more practical and convenient. Use the wildcard character asterisk (*) for any key properties that you do not need to match. The following table provides some examples of object names with wildcard key properties that match single or multiple MBeans.

Table 6. Examples object names with wildcard key properties. Include asterisks (*) to specify wildcard key properties.

:type=Server,	All MBeans of type Server
:node=Node1,type=Server,	All MBeans of type Server on Node1
:type=JVM,process=server1,node=Node1,	The JVM MBean in the server named server1 node Node1
:process=server1,	All MBeans in all servers named server1
:process=server1,node=Node1,	All MBeans in the server named server1 on Node1

You can locate an MBean by querying for it with object names that match key properties. The following example shows how to find the MBean for the node agent of node, MyNode:

```
String nodeName = "MyNode";
String query = "WebSphere:type=NodeAgent,node=" + nodeName + ",*";
ObjectName queryName = new ObjectName(query);
ObjectName nodeAgent = null;
Set s = adminClient.queryNames(queryName, null);
if (!s.isEmpty())
    nodeAgent = (ObjectName)s.iterator().next();
else
    System.out.println("Node agent MBean was not found");
```

- a. Build an ObjectName instance with a query string that specifies the key properties of type and node.

By using a wildcard for the remaining key properties, this pattern matches the object names for all MBeans of the type NodeAgent on the node MyNode. Because only one node agent per node exists, this information is sufficient to identify the MBean that you want.

- b. Give this ObjectName instance to the queryNames method of the AdminClient interface.

The AdminClient interface performs the remote call to the AdminService interface to obtain the set of MBean object names that match the query. The null second parameter to this method is a query expression (QueryExp) object that you can use as an additional query over the MBeans that match the ObjectName pattern in the first parameter.

- c. Use the set iterator to get the first and, in this case, only element.

The element is the MBean ObjectName instance of the node agent.

3. Use the MBean.

What a particular MBean can do depends on the management interface of that MBean. An MBean can declare:

- Attributes that you can obtain or set
- Operations that you can invoke
- Notifications for which you can register listeners

For the MBeans provided by WebSphere Application Server, you can find information about the interfaces they support in the MBean API documentation. The following example invokes one of the operations available on the NodeAgent MBean that you located previously. The following example starts the *MyServer* application server:

```
String opName = "launchProcess";
String signature[] = { "java.lang.String" };
String params[] = { "MyServer" };
try
{
    adminClient.invoke(nodeAgent, opName, params, signature);
}
catch (Exception e)
{
    System.out.println("Exception invoking launchProcess: " + e);
}
```

The `AdminClient.invoke` method is a generic means of invoking any operation on any MBean. The parameters are:

- The object name of the target MBean, `nodeAgent`
- The name of the operation, `opName`
- An object array that contains the operation parameters, `params`
- A string array that contains the operation signature, `signature`

The `launchProcess` operation in the example has a single parameter which is a string that identifies the server to start.

The `invoke` method returns an object instance, which the calling code can use to cast to the correct return type for the invoked operation. The `launchProcess` operation is declared void so that you can ignore the return value in this example.

4. Register for events.

In addition to managing resources, the JMX API also supports application monitoring for specific administrative events. Certain events produce notifications, for example, when a server starts. Administrative applications can register as listeners for these notifications. The WebSphere Application Server provides a full implementation of the JMX notification model, and provides additional function so you can receive notifications in a distributed environment. For a complete list of the notifications emitted from product MBeans, refer to the `com.ibm.websphere.management.NotificationConstants` class in the MBean API documentation.

The following example shows how an object can register for event notifications that are emitted from an MBean using the `ObjectName` node agent:

```
adminClient.addNotificationListener(nodeAgent, this, null, null);
```

In this example, the first parameter is the `ObjectName` for the node agent MBean. The second parameter identifies the listener object, which must implement the `NotificationListener` interface. In this case, the calling object is the listener. The third parameter is a filter that you can use to indicate which notifications you want to receive. When you leave this value as `null`, you receive all notifications from this MBean. The final parameter is a handback object that you can use to set the JMX API to return to you when it emits a notification.

If your MBean is located on another server in the cell, you can receive its notifications even though your administrative client program might be connected to the deployment manager server. All notifications flow to the upstream server. For example, a notification from an application server first flows to the local node agent and then to the deployment manager.

Another enhanced feature that Application Server provides is the ability to register as a notification listener of multiple MBeans with one call. This registration is done through the `addNotificationListenerExtended` method of the `AdminClient` interface, an extension of the standard JMX `addNotificationListener` method. This extension method even lets you register for MBeans that are not currently active. This registration is important in situations where you want to monitor events from resources that can be stopped and restarted during the lifetime of your administrative client program.

5. Handle the events.

Objects receive JMX event notifications through the `handleNotification` method, which is defined by the `NotificationListener` interface and which any event receiver must implement. The following example is an implementation of the `handleNotification` method that reports the notifications that it receives:

```
public void handleNotification(Notification n, Object handback)
{
    System.out.println("*****");
    System.out.println("* Notification received at " + new Date().toString());
    System.out.println("* type      = " + ntfyObj.getType());
    System.out.println("* message = " + ntfyObj.getMessage());
    System.out.println("* source  = " + ntfyObj.getSource());
    System.out.println(
        "* seqNum   = " + Long.toString(ntfyObj.getSequenceNumber()));
}
```



```

        System.out.println("* timeStamp = " + new Date(ntfyObj.getTimeStamp()));
        System.out.println("* userData = " + ntfyObj.getUserData());
        System.out.println("*****");
    }
}

```

Results

The administrative client can handle event notifications that are emitted from an MBean.

Note: If a client program registers a notification listener through an RMI or JSR160RMI connector and the ORB thread does not stop running, and thus prevents the Java virtual machine from exiting, add a `System.exit()` statement to the client program. The ORB starts a thread to handle notification propagation to the client. This thread does not automatically exit with the client main thread unless the main thread has a `System.exit()` statement. Place a `System.exit()` statement in a location in the client program that enables the ORB thread and main thread to stop processing. For example, place the `System.exit()` statement in a catch or finally clause of the client program main try block.

Example: Administrative client program

Copy the contents to a file named `AdminClientExample.java`. After changing the node name and server name to the appropriate values for your configuration, you can compile and run it using the instructions from [Creating a custom Java administrative client program using WebSphere Application Server administrative Java APIs](#)

```

import java.util.Date;
import java.util.Properties;
import java.util.Set;

import javax.management.InstanceNotFoundException;
import javax.management.MalformedObjectNameException;
import javax.management.Notification;
import javax.management.NotificationListener;
import javax.management.ObjectName;

import com.ibm.websphere.management.AdminClient;
import com.ibm.websphere.management.AdminClientFactory;
import com.ibm.websphere.management.exception.ConnectorException;

public class AdminClientExample implements NotificationListener
{
    private AdminClient adminClient;
    private ObjectName nodeAgent;
    private long ntfyCount = 0;

    public static void main(String[] args)
    {
        AdminClientExample ace = new AdminClientExample();

        // Create an AdminClient
        ace.createAdminClient();

        // Find a NodeAgent MBean
        ace.getNodeAgentMBean("ellington");

        // Invoke launchProcess
        ace.invokeLaunchProcess("server1");

        // Register for NodeAgent events
        ace.registerNotificationListener();

        // Run until interrupted
        ace.countNotifications();
    }
}

```

```

}

private void createAdminClient()
{
    // Set up a Properties object for the JMX connector attributes
    Properties connectProps = new Properties();
    connectProps.setProperty(
        AdminClient.CONNECTOR_TYPE, AdminClient.CONNECTOR_TYPE_SOAP);
    connectProps.setProperty(AdminClient.CONNECTOR_HOST, "localhost");
    connectProps.setProperty(AdminClient.CONNECTOR_PORT, "8879");

    // Get an AdminClient based on the connector properties
    try
    {
        adminClient = AdminClientFactory.createAdminClient(connectProps);
    }
    catch (ConnectorException e)
    {
        System.out.println("Exception creating admin client: " + e);
        System.exit(-1);
    }

    System.out.println("Connected to DeploymentManager");
}

private void getNodeAgentMBean(String nodeName)
{
    // Query for the ObjectName of the NodeAgent MBean on the given node
    try
    {
        String query = "WebSphere:type=NodeAgent,node=" + nodeName + ",*";
        ObjectName queryName = new ObjectName(query);
        Set s = adminClient.queryNames(queryName, null);
        if (!s.isEmpty())
            nodeAgent = (ObjectName)s.iterator().next();
        else
        {
            System.out.println("Node agent MBean was not found");
            System.exit(-1);
        }
    }
    catch (MalformedObjectNameException e)
    {
        System.out.println(e);
        System.exit(-1);
    }
    catch (ConnectorException e)
    {
        System.out.println(e);
        System.exit(-1);
    }

    System.out.println("Found NodeAgent MBean for node " + nodeName);
}

private void invokeLaunchProcess(String serverName)
{
    // Use the launchProcess operation on the NodeAgent MBean to start
    // the given server
    String opName = "launchProcess";
    String signature[] = { "java.lang.String" };
    String params[] = { serverName };
    boolean launched = false;
    try
    {
        Boolean b = (Boolean)adminClient.invoke(

```

```

nodeAgent, opName, params, signature);
    launched = b.booleanValue();
    if (launched)
        System.out.println(serverName + " was launched");
    else
        System.out.println(serverName + " was not launched");

}
catch (Exception e)
{
    System.out.println("Exception invoking launchProcess: " + e);
}
}

private void registerNotificationListener()
{
    // Register this object as a listener for notifications from the
    // NodeAgent MBean. Don't use a filter and don't use a handback
    // object.
    try
    {
        adminClient.addNotificationListener(nodeAgent, this, null, null);
        System.out.println("Registered for event notifications");
    }
    catch (InstanceNotFoundException e)
    {
        System.out.println(e);
        e.printStackTrace();
    }
    catch (ConnectorException e)
    {
        System.out.println(e);
        e.printStackTrace();
    }
}

public void handleNotification(Notification ntfyObj, Object handback)
{
    // Each notification that the NodeAgent MBean generates will result in
    // this method being called
    ntfyCount++;
    System.out.println("*****");
    System.out.println("* Notification received at " + new Date().toString());
    System.out.println("* type      = " + ntfyObj.getType());
    System.out.println("* message   = " + ntfyObj.getMessage());
    System.out.println("* source    = " + ntfyObj.getSource());
    System.out.println(
        "* seqNum   = " + Long.toString(ntfyObj.getSequenceNumber());
    System.out.println("* timeStamp = " + new Date(ntfyObj.getTimeStamp());
    System.out.println("* userData  = " + ntfyObj.getUserData());
    System.out.println("*****");
}

private void countNotifications()
{
    // Run until killed
    try
    {
        while (true)
        {
            Thread.currentThread().sleep(60000);
            System.out.println(ntfyCount + " notification have been received");
        }
    }
    catch (InterruptedException e)

```

```

    {
    }
}

```

Example: Administrative client program:

This example is a complete administrative client program.

Copy the contents to a file named `AdminClientExample.java`. After changing the node name and server name to the appropriate values for your configuration, you can compile and run it using the instructions from [Creating a custom Java administrative client program using WebSphere Application Server administrative Java APIs](#)

```

import java.util.Date;
import java.util.Properties;
import java.util.Set;

import javax.management.InstanceNotFoundException;
import javax.management.MalformedObjectNameException;
import javax.management.Notification;
import javax.management.NotificationListener;
import javax.management.ObjectName;

import com.ibm.websphere.management.AdminClient;
import com.ibm.websphere.management.AdminClientFactory;
import com.ibm.websphere.management.exception.ConnectorException;

public class AdminClientExample implements NotificationListener
{
    private AdminClient adminClient;
    private ObjectName nodeAgent;
    private long ntfyCount = 0;

    public static void main(String[] args)
    {
        AdminClientExample ace = new AdminClientExample();

        // Create an AdminClient
        ace.createAdminClient();

        // Find a NodeAgent MBean
        ace.getNodeAgentMBean("ellington");

        // Invoke launchProcess
        ace.invokeLaunchProcess("server1");

        // Register for NodeAgent events
        ace.registerNotificationListener();

        // Run until interrupted
        ace.countNotifications();
    }

    private void createAdminClient()
    {
        // Set up a Properties object for the JMX connector attributes
        Properties connectProps = new Properties();
        connectProps.setProperty(
            AdminClient.CONNECTOR_TYPE, AdminClient.CONNECTOR_TYPE_SOAP);
        connectProps.setProperty(AdminClient.CONNECTOR_HOST, "localhost");
        connectProps.setProperty(AdminClient.CONNECTOR_PORT, "8879");

        // Get an AdminClient based on the connector properties
    }
}

```

```

    try
    {
        adminClient = AdminClientFactory.createAdminClient(connectProps);
    }
    catch (ConnectorException e)
    {
        System.out.println("Exception creating admin client: " + e);
        System.exit(-1);
    }

    System.out.println("Connected to DeploymentManager");
}

private void getNodeAgentMBean(String nodeName)
{
    // Query for the ObjectName of the NodeAgent MBean on the given node
    try
    {
        String query = "WebSphere:type=NodeAgent,node=" + nodeName + ",*";
        ObjectName queryName = new ObjectName(query);
        Set s = adminClient.queryNames(queryName, null);
        if (!s.isEmpty())
            nodeAgent = (ObjectName)s.iterator().next();
        else
        {
            System.out.println("Node agent MBean was not found");
            System.exit(-1);
        }
    }
    catch (MalformedObjectNameException e)
    {
        System.out.println(e);
        System.exit(-1);
    }
    catch (ConnectorException e)
    {
        System.out.println(e);
        System.exit(-1);
    }

    System.out.println("Found NodeAgent MBean for node " + nodeName);
}

private void invokeLaunchProcess(String serverName)
{
    // Use the launchProcess operation on the NodeAgent MBean to start
    // the given server
    String opName = "launchProcess";
    String signature[] = { "java.lang.String" };
    String params[] = { serverName };
    boolean launched = false;
    try
    {
        Boolean b = (Boolean)adminClient.invoke(
nodeAgent, opName, params, signature);
        launched = b.booleanValue();
        if (launched)
            System.out.println(serverName + " was launched");
        else
            System.out.println(serverName + " was not launched");
    }
    catch (Exception e)
    {
        System.out.println("Exception invoking launchProcess: " + e);
    }
}

```

```

    }
}

private void registerNotificationListener()
{
    // Register this object as a listener for notifications from the
    // NodeAgent MBean. Don't use a filter and don't use a handback
    // object.
    try
    {
        adminClient.addNotificationListener(nodeAgent, this, null, null);
        System.out.println("Registered for event notifications");
    }
    catch (InstanceNotFoundException e)
    {
        System.out.println(e);
    }
    catch (ConnectorException e)
    {
        System.out.println(e);
    }
}

public void handleNotification(Notification ntfyObj, Object handback)
{
    // Each notification that the NodeAgent MBean generates will result in
    // this method being called
    ntfyCount++;
    System.out.println("*****");
    System.out.println("* Notification received at " + new Date().toString());
    System.out.println("* type      = " + ntfyObj.getType());
    System.out.println("* message   = " + ntfyObj.getMessage());
    System.out.println("* source    = " + ntfyObj.getSource());
    System.out.println(
        "* seqNum    = " + Long.toString(ntfyObj.getSequenceNumber());
    System.out.println("* timeStamp = " + new Date(ntfyObj.getTimeStamp());
    System.out.println("* userData  = " + ntfyObj.getUserData());
    System.out.println("*****");
}

private void countNotifications()
{
    // Run until killed
    try
    {
        while (true)
        {
            Thread.currentThread().sleep(60000);
            System.out.println(ntfyCount + " notification have been received");
        }
    }
    catch (InterruptedException e)
    {
    }
}
}

```

Creating a Java Management Extensions client program using the Java Management Extensions Remote application programming interface

This topic describes how to develop and build a Java Management Extensions (JMX) client program that is compliant with JMX Remote application programming interface (JSR 160). After you have a working JMX client program, you can use it to manage WebSphere Application Server or non-WebSphere Application Server systems.

Before you begin

This task assumes a basic familiarity with JSR 160 and JMX application programming interface (API) programming. For information on JSR 160, see <http://www.jcp.org/en/jsr/detail?id=160>. For information on the Java APIs, view the application programming interfaces documentation.

About this task

When you develop and run JMX clients that use various JMX connectors and that have security enabled, use the following guidelines. When you follow these guidelines, you guarantee the behavior among different implementations of JMX connectors. Any programming model that strays from these guidelines is unsupported.

1. Create and use a single JMX client before you create and use another JMX client.
2. Create and use a JMX client on the same thread.
3. Use one of the following ways to specify a user ID and password to create a new JMX client:
 - Specify a default user ID and password in the property file.
 - Specify a user ID and password other than the default. After you create a JMX client with a nondefault user ID and password, specify the nondefault user ID and password when you create subsequent JMX clients.

Procedure

1. Develop a JMX client program.
2. Build the JMX remote client program.

Compile the program with the `javac` command and provide the location of the `ibm.admin.thinclient.jar` file in the classpath argument.

For example, if your `ibm.admin.thinclient.jar` file is in the `/opt/resources/ibm.admin.thinclient.jar` path, and you want to compile the `JMXRemoteClientApp.java` file in the current directory, use the following settings and commands:

```
CLASSPATH=/opt/resources/ibm.ws.admin.thinclient.jar:${CLASSPATH}
export CLASSPATH
${JAVA_HOME}/bin/javac JMXRemoteClientApp.java
```

3. Run the JMX client program.

Run the JMX client program by setting up the runtime environment so that the program can find all of the prerequisites. Many of the batch or script files in the `bin` directory under the installation root perform a similar function. The following example is a batch file that runs the `JMXRemoteClientApp` JMX client program:

```
#!/bin/sh
CONNECTORPROPS=-Dcom.ibm.CORBA.configURL=<location of sas.client.props>
JAVA_HOME=<location of Java>
TAC_CLASSPATH=<location of ibm.admin.thinclient.jar>
```

```
"${JAVA_HOME}/bin/java" \
-Djava.ext.dirs="${JAVA_HOME}/jre/lib/ext" \
-classpath "${TAC_CLASSPATH}" ${CONNECTORPROPS} JMXRemoteClientApp $@
```

(The contents of the previous batch file are split on multiple lines for publication.)

Results

You have developed, built, and run a JMX client program that is JSR 160 compliant.

Developing a Java Management Extensions client program using Java Management Extensions Remote application programming interface

This topic describes how to develop a Java Management Extensions (JMX) connector specification and JMX Remote application programming interface (API) (JSR 160). The program can communicate by Remote Method Invocation over Internet Inter-ORB Protocol (RMI-IIOP)

Before you begin

This topic assumes a basic understanding of JSR 160, JMX APIs, and managed beans (MBeans). For more information on JSR 160, see the JSR 160: Java Management Extensions (JMX) Remote API at <http://www.jcp.org/en/jsr/detail?id=160>. For more information on the JMX APIs and on MBeans, view the application programming interfaces documentation.

About this task

You can administer your WebSphere Application Server environment through the administrative console, the wsadmin utility, or Java Management Extensions (JMX) programming. Complete this task to develop a JMX remote client program using the JMX remote API so that you can administer your environment through JMX programming.

Procedure

1. Specify the JMX connector address for the server through the JMXServiceURL class.

The value of the JMX service URL is:

```
service:jmx:rmi://" + host + ":" + port + "/jndi/JMXConnector"
```

For example, if the target server host is sales.xyz.com and the listening port is 1234, the JMX service URL is:

```
service:jmx:rmi://sales.xyz.com:1234/jndi/JMXConnector
```

You can find the value for *port* in the Ports table of the console server settings page or in the serverindex.xml file that includes the target server. If the URL does not specify a value for *host*, the product uses the default value of localhost. If the URL does not specify a value for *port*, the product uses the default value of 2809.

When connecting to an administrative agent, add the administrative agent JMX connector port number to the end of the URL. For example, if the administrative agent JMX connector host is sales.xyz.com and the port is 6789, then use the following URL:

```
service:jmx:rmi://sales.xyz.com:6789/jndi/JMXConnector6789
```

2. Set the Java Naming and Directory Interface (JNDI) provider URL property to use the administrative name service for the product.

The JNDI provider URL property is javax.naming.Context.PROVIDER_URL. The administrative name service is WsnAdminNameService.

3. If the client uses security, set the -Dcom.ibm.CORBA.ConfigURL and -Dcom.ibm.SSL.ConfigURL system properties in the client Java virtual machine (JVM).

Without the -Dcom.ibm.CORBA.ConfigURL and -Dcom.ibm.SSL.ConfigURL system properties set to valid system properties files, the client does not work properly when security is enabled. The recommended way to run the JMX connector client is as an administrative thin client.

- -Dcom.ibm.CORBA.ConfigURL=file:app_client_root/properties/sas.client.props
- -Dcom.ibm.SSL.ConfigURL=file:app_client_root/properties/ssl.client.props

Typically, you can copy the properties files from an installation profile directory, preferably from the target server profile directory.

4. Specify the user ID and password for the server, if security is enabled.
5. Establish the JMX connection.
6. Get the MBean server connection instance.

Results

You have established a connection to the deployment manager through an RMI connection and started the application server through the node agent MBean.

Example

Use the following thin client code example to create and use the JMX client.

Some statements are split on multiple lines for printing purposes.

```
import java.io.File;
import java.util.Date;
import java.util.Set;
import java.util.Hashtable;

import javax.management.Notification;
import javax.management.NotificationListener;
import javax.management.ObjectName;
import javax.management.MBeanServerConnection;
import javax.management.remote.JMXConnector;
import javax.management.remote.JMXConnectorFactory;
import javax.management.remote.JMXServiceURL;

public class JMXRemoteClientApp implements NotificationListener {

    private MBeanServerConnection mbsc = null;
    private ObjectName nodeAgent;
    private ObjectName jvm;
    private long ntfyCount = 0;
    private static String userid = null;
    private static String pwd = null;

    public static void main(String[] args)
    {
        try {

            JMXRemoteClientApp client = new JMXRemoteClientApp();

            String host=args[0];
            String port=args[1];
            String nodeName =args[2];
            userid =args[3];
            pwd = args[4];

            client.connect(host,port);

            // Find a node agent MBean
            client.getNodeAgentMBean(nodeName);

            // Invoke the launch process.
            client.invokeLaunchProcess("server1");

            // Register for node agent events
            client.registerNotificationListener();

            // Run until interrupted.
            client.countNotifications();

        } catch (Exception e) {
            e.printStackTrace();
        }
    }

    private void connect(String host,String port) throws Exception
    {
        String jndiPath="/WsnAdminNameService#JMXConnector";

        JMXServiceURL url =
            new JMXServiceURL("service:jmx:iiop://"+host+"/jndi/corbaname:iiop:"+host+": "+port+jndiPath);

        Hashtable h = new Hashtable();

        //Specify the user ID and password for the server if security is enabled on server.
```

```

System.out.println("userid is " + userid);
System.out.println("password is " + pwd);
if ((userid.length() != 0) && (pwd.length() != 0)) {
    System.out.println("adding userid and password to credentials...");
    String[] credentials = new String[] {userid , pwd };
    h.put("jmx.remote.credentials", credentials);
} else {
    System.out.println("No credentials provided.");
}

//Establish the JMX connection.

JMXConnector jmxc = JMXConnectorFactory.connect(url, h);

//Get the MBean server connection instance.

mbsc = jmxc.getMBeanServerConnection();

System.out.println("Connected to DeploymentManager");
}

private void getNodeAgentMBean(String nodeName)
{
    // Query for the object name of the node agent MBean on the given node
    try {
        String query = "WebSphere:type=NodeAgent,node=" + nodeName + ",*";
        ObjectName queryName = new ObjectName(query);
        Set s = mbsc.queryNames(queryName, null);
        if (!s.isEmpty()) {
            nodeAgent = (ObjectName)s.iterator().next();
            System.out.println("NodeAgent mbean found "+ nodeAgent.toString());
        } else {
            System.out.println("Node agent MBean was not found");
            System.exit(-1);
        }
    } catch (Exception e) {
        System.out.println(e);
        System.exit(-1);
    }
}

private void invokeLaunchProcess(String serverName)
{
    // Use the launch process on the node agent MBean to start
    // the given server.
    String opName = "launchProcess";
    String signature[] = { "java.lang.String" };
    String params[] = { serverName };
    boolean launched = false;
    try {
        Boolean b = (Boolean)mbsc.invoke(nodeAgent, opName, params, signature);
        launched = b.booleanValue();
        if (launched)
            System.out.println(serverName + " was launched");
        else
            System.out.println(serverName + " was not launched");
    } catch (Exception e) {
        System.out.println("Exception invoking launchProcess: " + e);
    }
}

private void registerNotificationListener()
{
    // Register this object as a listener for notifications from the
    // node agent MBean. Do not use a filter and do not use a handback
    // object.
    try {
        mbsc.addNotificationListener(nodeAgent, this, null, null);
        System.out.println("Registered for event notifications");
    } catch (Exception e) {
        System.out.println(e);
    }
}

public void handleNotification(Notification ntfyObj, Object handback)
{
    // Each notification that the node agent MBean generates results in
    // a call to this method.
    ntfyCount++;
    System.out.println("*****");
    System.out.println("* Notification received at " + new Date().toString());
    System.out.println("* type      = " + ntfyObj.getType());
    System.out.println("* message = " + ntfyObj.getMessage());
    System.out.println("* source  = " + ntfyObj.getSource());
}

```

```

System.out.println(
    "* seqNum    = " + Long.toString(ntfyObj.getSequenceNumber());
System.out.println("* timeStamp = " + new Date(ntfyObj.getTimeStamp());
System.out.println("* userData  = " + ntfyObj.getUserData());
System.out.println("*****");
}

private void countNotifications()
{
    // Run until stopped.
    try {
        while (true) {
            Thread.currentThread().sleep(60000);
            System.out.println(ntfyCount + " notification have been received");
        }
    } catch (InterruptedException e) {
    }
}
}

```

Extending the WebSphere Application Server administrative system with custom MBeans

You can extend the WebSphere Application Server administration system by supplying and registering new Java Management Extensions (JMX) MBeans (see JMX 1.x Specification for details) in one of the WebSphere processes.

About this task

JMX MBeans represent the management interface for a particular piece of logic. All of the managed resources within the standard product infrastructure are represented as JMX MBeans. There are a variety of ways in which you can create your own MBeans and register them with the JMX MBeanServer running in any WebSphere process. For more information, see MBean Java application programming interface (API) documentation.

Procedure

1. Create custom JMX MBeans.

You have some alternatives to select from, when creating MBeans to extend the product administrative system. You can use any existing JMX MBean from another application. You can register any MBean that you tested in a JMX MBean server outside of the WebSphere Application Server environment in a product process, including standard MBeans, dynamic MBeans, open MBeans, and model MBeans.

In addition to any existing JMX MBeans, and ones that were written and tested outside of the product environment, you can use the special distributed extensions provided by WebSphere and create a WebSphere ExtensionMBean provider. This alternative provides better integration with all of the distributed functions of the product administrative system. An ExtensionMBean provider implies that you supply an XML file that contains an MBean Descriptor based on the DTD shipped with the product. This descriptor tells the WebSphere system all of the attributes, operations, and notifications that your MBean supports. With this information, the WebSphere system can route remote requests to your MBean and register remote Listeners to receive your MBean event notifications.

All of the internal WebSphere MBeans follow the Model MBean pattern. Pure Java classes supply the real logic for management functions, and the WebSphere MBeanFactory class reads the description of these functions from the XML MBean Descriptor and creates an instance of a ModelMBean that matches the descriptor. This ModelMBean instance is bound to your Java classes and registered with the MBeanServer running in the same process as your classes. Your Java code now becomes callable from any WebSphere Application Server administrative client through the ModelMBean created and registered to represent it.

User MBeans that run on both the WebSphere Application Server distributed platforms and the WebSphere Application Server for z/OS platform may require special coding to function properly in the z/OS multiprocess model. On distributed platforms where each application server runs in a single Java virtual machine (JVM), there is only one MBean server. The MBean server controls all MBeans that are registered within that application server. On the z/OS platform, there is a control process and a

federation of servant processes, each with their own MBean server. The control process has its own MBean proxy server to distribute requests among the servant processes. See the detailed discussion of the JMX MBean multiprocess model request flow.

2. Optionally define an explicit MBean security policy.

If you do not define an MBean security policy, the product uses the default security policy.

3. Register the new MBeans. There are various ways to register your MBean.

You can register your MBean with the WebSphere Application Server administrative service.

You can register your MBean with the MBeanServer in a WebSphere Application Server process. The following list describes the available options in order of preference:

- Go through the MBeanFactory class. If you want the greatest possible integration with the WebSphere Application Server system, then use the MBeanFactory class to manage the life cycle of your MBean through the activateMBean and deactivateMBean methods of the MBeanFactory class. Use these methods, by supplying a subclass of the RuntimeCollaborator abstract superclass and an XML MBean descriptor file. Using this approach, you supply a pure Java class that implements the management interface defined in the MBean descriptor. The MBeanFactory class creates the actual ModelMBean and registers it with the product administrative system on your behalf.

This option is recommended for registering model MBeans.

- Use the JMXManageable and CustomService interface. You can make the process of integrating with WebSphere administration even easier by implementing a CustomService interface that also implements the JMXManageable interface. Using this approach, you can avoid supplying the RuntimeCollaborator. When your CustomService interface is initialized, the WebSphere MBeanFactory class reads your XML MBean descriptor file and creates, binds, and registers an MBean to your CustomService interface automatically. After the shutdown method of your CustomService is called, the product system automatically deactivates your MBean.
- Go through the AdminService interface. You can call the registerMBean() method on the AdminService interface and the invocation is delegated to the underlying MBeanServer for the process, after appropriate security checks. You can obtain a reference to the AdminService using the getAdminService() method of the AdminServiceFactory class.

This option is recommended for registering standard, dynamic, and open MBeans. Implement the UserCollaborator class to use the MBeans and to provide a consistent level of support for them across distributed and z/OS platforms.

For the z/OS platform, an MBean registered through the registerMBean() method on the AdminService interface is not visible from outside the server and can only be invoked from within the servant process in which it was registered.

- Get MBeanServer instances directly. You can get a direct reference to the JMX MBeanServer instance running in any product process, by calling the getMBeanServer() method of the MBeanFactory class. You get a reference to the MBeanFactory class by calling the getMBeanFactory() method of the AdminService interface.

When a custom MBean is registered directly with the MBean server, the MBean object name is enhanced with the cell, node and process name keys by default. This registration allows the MBean to participate in the distributed features of the administrative system. You can turn off the default behavior by setting the `com.ibm.websphere.mbeans.disableRouting` custom property.

See the *Installing your application serving environment* PDF for more information on the `com.ibm.websphere.mbeans.disableRouting` custom property.

Results

Regardless of the approach used to create and register your MBean, you must set up proper Java 2 security permissions for your new MBean code. The WebSphere AdminService and MBeanServer are tightly protected using Java 2 security permissions and if you do not explicitly grant your code base permissions, security exceptions are thrown when you attempt to invoke methods of these classes. If you are supplying your MBean as part of your application, you can set the permissions in the `was.policy` file that you supply as part of your application metadata. If you are using a CustomService interface or other

code that is not delivered as an application, you can edit the `library.policy` file in the node configuration, or even the `server.policy` file in the `properties` directory for a specific installation.

Example: The `SampleStateMBean` MBean

MBeanDescriptor

```
<?xml version="1.0" encoding="UTF-8"?>
<!DOCTYPE MBean SYSTEM "MbeanDescriptor.dtd">
<MBean type="SampleStateMBean"
  aggregationHandlerClass="com.ibm.ws390.sample.SampleStateAggregationHandler"
  eventHandlerClass="com.ibm.ws390.sample.SampleStateEventHandler"
  invocationHandlerClass="com.ibm.ws390.sample.SampleStateInvocationHandler"
  stateObjectClass="com.ibm.ws390.sample.SampleState"
  version="6.0"
  platform="dynamicproxy"
  description="Sample State MBean for the documentation example.">

  <attribute description="The name of the MBean."
    getMethod="getMBeanName" name="mbeanName" type="java.lang.String"
    proxyInvokeType="unicall"/>
  <attribute description="The state of the MBean." name="state"
    getMethod="getState" setMethod="setState" type="java.lang.String"
    proxyInvokeType="multicall" proxySetterInvokeType="multicall"/>
  <operation
    description="Initialize the State MBean."
    impact="ACTION" name="initializeState" role="operation"
    targetObjectType="objectReference" type="void" proxyInvokeType="multicall">
    <signature>
      <parameter description="The name of the MBean."
        name="mbeanName" type="java.lang.String"/>
      <parameter description="The initial state of the MBean."
        name="mbeanName" type="java.lang.String"/>
    </signature>
  </operation>

  <notification name="j2ee.state.starting" severity="6" log="false"
    description="This sample state MBean is in starting state.">
    <notificationType>j2ee.state.starting</notificationType>
  </notification>
  <notification name="j2ee.state.running" severity="6" log="false"
    description="This sample state MBean is in running state.">
    <notificationType>j2ee.state.running</notificationType>
  </notification>
  <notification name="j2ee.state.stopping" severity="6" log="false"
    description="This sample state MBean is in stopping state.">
    <notificationType>j2ee.state.stopping</notificationType>
  </notification>
  <notification name="j2ee.state.stopped" severity="6" log="false"
    description="This sample state MBean is in stopped state.">
    <notificationType>j2ee.state.stopped</notificationType>
  </notification>
</MBean>
```

SampleState implementation

```
package com.ibm.ws390.sample;

import com.ibm.ejs.ras.Tr;
import com.ibm.ejs.ras.TraceComponent;
import java.io.Serializable;
import com.ibm.websphere.management.dynamicproxy.StateObject;

public class SampleState extends StateObject {

    private static TraceComponent tc =
        Tr.register(SampleState.class, "SampleState", null);
```

```

// Package protected STATE constants.
static final String STATE_STARTING = "j2ee.state.starting";
static final String STATE_RUNNING = "j2ee.state.running";
static final String STATE_STOPPING = "j2ee.state.stopping";
static final String STATE_STOPPED = "j2ee.state.stopped";

// Dynamicproxy State is initialized with STOPPED state.
private String state = STATE_STOPPED;

public SampleState() {
    if (tc.isEntryEnabled()) Tr.entry(tc,"<init>");

    // State is initialized during "state" initialization above,
    // but can also be initialized here in the constructor as well.
    /*
    state = "WebSphere Application Server for z/OS ready for e-business";
    */

    if (tc.isEntryEnabled()) Tr.exit(tc,"<init>");
}

public synchronized String getState() {
    if (tc.isEntryEnabled()) Tr.entry(tc,"getState");
    if (tc.isEntryEnabled()) Tr.exit(tc,"getState",state);
    return state;
}

public synchronized void setState(String state) {
    if (tc.isEntryEnabled()) Tr.entry(tc,"setState",state);
    this.state = state;
    if (tc.isEntryEnabled()) Tr.exit(tc,"setState");
}

public synchronized String getStateObjectInfo() {
    return state;
}
}

```

SampleStateAggregationHandler implementation

```

package com.ibm.ws390.sample;

import com.ibm.websphere.management.dynamicproxy.AggregationHandler;
import com.ibm.websphere.management.dynamicproxy.StateObject;

import com.ibm.ejs.ras.Tr;
import com.ibm.ejs.ras.TraceComponent;

public class SampleStateAggregationHandler implements AggregationHandler {

    private static TraceComponent tc =
        Tr.register(SampleStateAggregationHandler.class,"SampleState",null);

    /**
     * Return an aggregated result from a multicall Mbean operation which
     * compiles through all servant MBeans' results and returns a respective
     * single return value for an invoked method.
     *
     * @param methodName      MBean method name
     * @param params          MBean method parameters
     * @param signatures      MBean method signatures
     * @param servantMBeanResults Result of each servant MBean instances
     *                          invoked by the dynamicproxy multicast
     *                          invocation.
     *
     * Note: this value can be "null" OR can be
     * an array of "null"s in case return value
     */

```

```

*           of the method is "void." Implementation
*           of this method MUST handle this case to
*           avoid a <code>NullPointerException</code>.
* @param stateObject
*         MBean provider provided <code>StateObject</code> used by
*         dynamicproxy MBean in CR to manage its state. Note: this object
*         MAY BE null if "stateObjectClass" was not specified OR internal
*         error occurred during initialization of this dynamicproxy MBean.
*         Implementation MUST properly handle "null" input.
*
* @return aggregated result as defined by MBean xml for specified
*         MBean operation.
*/
public Object aggregateResults(String methodName,
                              Object[] params,
                              String[] signatures,
                              Object[] servantMBeanResults,
                              StateObject stateObject) {

    if (tc.isEntryEnabled()) Tr.entry(tc,"aggregateResults",methodName);

    // As you can see from the MBeanDescriptor of SampleStateMBean,
    // it declares the following four methods:
    // 1. String getMBeanName()           [proxyInvokeType == unicast]
    // 2. String getState()               [proxyInvokeType == multicast]
    // 3. void setState(String)          [proxyInvokeType == multicast]
    // 4. void initializeState()         [proxyInvokeType == multicast]
    //
    // Looking at the above methods, only method that requires aggregation
    // is #2 getState method which is a multicast MBean operation AND
    // it returns a value that can be aggregated.
    //
    // In this example, we simply take each servants' getState MBean
    // request result and concatenate them into one long String that
    // displays each servants' state.
    if (methodName.equals("getState")) {
        StringBuffer stateBuf = new StringBuffer();

        for (int i=0; i<servantMBeanResults.length; i++) {
            stateBuf.append("SERVANT #" + i + " state ==|" +
                servantMBeanResults[i] + "|== ");
        }
        return stateBuf.toString();
    }
    // If we also had an example method which returns say an int,
    // getNumberOfMBeans(), it can take the similar approach
    // and to add each servants' getNumberOfMBeans() result together here.
    /* example added for non-existent method: int getNumberOfMBeans()
    else if (methodName.equals("getNumberOfMBeans")) {
        int aggregatedResult = 0;

        for (int i=0; i<servantMBeanResults.length; i++) {
            aggregatedResult += (int) servantMBeanResults[i];
        }
        return aggregatedResult;
    }
    */

    return methodName + " is NOT handled by " + getClass().getName() + "!";
}
}

```

Best practices for standard, dynamic, and open MBeans

This topic discusses recommended guidelines for standard, dynamic, and open MBeans.

The underlying interface for the WebSphere Application Server administrative service is AdminService. Remote access occurs through the AdminControl scripting object.

The product provides a special runtime collaborator that you use with standard, dynamic or open custom MBeans to register the custom MBeans with the WebSphere Application Server administrative service. The standard, dynamic, and open MBeans display in the administrative service as model MBeans. The administrative service uses the capabilities available to MBeans that are registered with the administrative service.

The MBean registration and capabilities are as follows:

Table 7. MBean registration and capabilities. Examine the registration and capabilities for an MBean type.

MBean type	Registered with:	Capabilities
Model, and optionally standard, dynamic, or open	WebSphere Application Server administrative service	Local access is through the WebSphere Application Server administrative service or the MBean server. Remote access is through the WebSphere Application Server administrative service, and WebSphere Application Server security. Remote access is also through z/OS system extensions.
Standard, dynamic, or open	MBean server	Local access is through the WebSphere Application Server administrative service or the MBean server on the distributed platform. Local access is only through the MBean server. Remote access is through the WebSphere Application Server administrative service, the Java Management Extensions (JMX) Remote application programming interface (API) (JSR 160) client code, and WebSphere Application Server security.

Creating and registering standard, dynamic, and open custom MBeans

You can create standard, dynamic, and open custom MBeans and register them with the product administrative service.

Before you begin

This task assumes a basic familiarity with MBean programming. For information on MBean programming, see MBean Java application programming interface (API) documentation.

Do not define new classes as parameters for your MBeans. The classes might not be available in all processes. If you must define a new class, ensure that the class is available on all processes, including the deployment manager, the node agents, and the application servers. If the class is not available for a process, the ClassNotFoundException exception occurs for the new class when you attempt to access it.

About this task

Perform the following tasks to create and register a standard, dynamic, or open custom MBean.

Procedure

1. Create your particular MBean class or classes.
2. Write an MBean descriptor in the XML language for your MBean.
3. Register your MBean by inserting code that uses the WebSphere Application Server runtime `com.ibm.websphere.management.UserMBeanCollaborator` collaborator class into your application code.
4. Package the class files for your MBean interface and implementation, the descriptor XML file, and your application Java archive (JAR) file.

Results

After you successfully complete the steps, you have a standard, dynamic, or open custom MBean that is registered and activated with the product administrative service.

Example

The following example shows how to create and register a standard MBean with the administrative service:

SnoopMBean.java:

```
/**
 * Use the SnoopMBean MBean, which has a standard mbean interface.
 */
public interface SnoopMBean {
    public String getIdentification();
    public void snoopy(String parm1);
}
```

SnoopMBeanImpl.java:

```
/**
 * SnoopMBeanImpl - SnoopMBean implementation
 */
public class SnoopMBeanImpl implements SnoopMBean {
    public String getIdentification() {
        System.out.println(">>> getIdentification() called...");
        return "snoopy!";
    }

    public void snoopy(String parm1) {
        System.out.println(">>> snoopy(" + parm1 + ") called...");
    }
}
```

Define the following MBean descriptor for your MBean in an .xml file. The getIdentification method is set to run with the unicast option and the snoopy method is set to use the multicast option. These options are used only for z/OS platform applications. The WebSphere Application Server for z/OS options are not applicable to the distributed platforms, but they do not need to be removed. The options are ignored on the distributed platforms. . Some statements are split on multiple lines for printing purposes.

gotcha: If you are running in a multiple JVM environment you must include the type property in the MBean descriptor.

SnoopMBean.xml:

```
<?xml version="1.0" encoding="UTF-8"?>
<!DOCTYPE MBean SYSTEM "MbeanDescriptor.dtd">
<MBean type="SnoopMBean"
  version="5.0"
  platform="dynamicproxy"
  description="Sample SnoopMBean to be initialized inside an EJB.">

  <attribute name="identification" getMethod="getIdentification"
  type="java.lang.String" proxyInvokeType="unicast"/>

  <operation name="snoopy" role="operation" type="void" targetObjectType="objectReference"
  impact="ACTION" proxyInvokeType="multicast">
    <signature>
```

```

    <parameter name="parm1" description="test parameter" type="java.lang.String"/>
  </signature>
</operation>
</MBean>

```

Assume that your MBean is used in an enterprise bean. Register your MBean in the enterprise bean `ejbCreate` method and unregister it in the `ejbRemove` method.

```

//The method MBeanFactory.activateMBean() requires four parameters:
//String type: The type value that you put in this MBean's descriptor. For this example
//the string type is SnoopMBean.
//RuntimeCollaborator co: The UserMBeanCollaborator user MBean collaborator instance
//that you create
//String id: Unique name that you pick
//String descriptor: The full path to the MBean descriptor file

import com.ibm.websphere.management.UserMBeanCollaborator;
//Import other classes here.
.
.
.
static private ObjectName snoopyON = null;
static private Object lockObj = "this is a lock";
.
.
.
/**
 * ejbCreate method: Register your Mbean.
 */
public void ejbCreate() throws javax.ejb.CreateException {
    synchronized (lockObj) {
        System.out.println(">>> SnoopMBean activating for --|" + this + "|--");
        if (snoopyON != null) {
            return;
        }
        try {
            System.out.println(">>> SnoopMBean activating...");
            MBeanFactory mbeanFactory = AdminServiceFactory.getMBeanFactory();
            RuntimeCollaborator snoop = new UserMBeanCollaborator(new SnoopMBeanImpl());
            snoopyON = mbeanFactory.activateMBean("SnoopMBean", snoop, "snoopMBeanId",
"SnoopMBean.xml");
            System.out.println(">>> SnoopMBean activation COMPLETED! --|" + snoopyON + "|--");
        } catch (Exception e) {
            System.out.println(">>> SnoopMBean activation FAILED:");
            e.printStackTrace();
        }
    }
}
.
.
.
/**
 * ejbRemove method: Unregister your MBean.
 */
public void ejbRemove() {
    synchronized (lockObj) {
        System.out.println(">>> SnoopMBean Deactivating for --|" + this + "|--");
        if (snoopyON == null) {
            return;
        }
        try {
            System.out.println(">>> SnoopMBean Deactivating ==|" + snoopyON + "|== for --|"
+ this + "|--");
            MBeanFactory mbeanFactory = AdminServiceFactory.getMBeanFactory();
            mbeanFactory.deactivateMBean(snoopyON);
            System.out.println(">>> SnoopMBean Deactivation COMPLETED!");
        }
    }
}

```

```

        } catch (Exception e) {
            System.out.println(">>> SnoopMBean Deactivation FAILED:");
            e.printStackTrace();
        }
    }
}

```

What to do next

Compile the MBean Java files and package the resulting class files with the descriptor .xml file, into the enterprise bean JAR file.

Setting Java 2 security permissions

You must configure Java 2 security permissions to use Java Management Extension and WebSphere Application Server administrative methods.

Before you begin

When you enable Java 2 security, you must grant Java 2 security permissions to application-specific code for Java Management Extensions (JMX) and WebSphere Application Server administrative privileges. With these permissions, your application code can call WebSphere Application Server administrative methods and JMX methods.

About this task

If you are using Java 2 security then you need to verify that your extensions and application server can access the required resources. The following steps show how to configure access for JMX and the application server administrative methods.

Procedure

- Use the following permission to invoke all the JMX class methods and interface methods:

```
permission javax.management.MBeanPermission "*", "*";
```

Consult the application programming interfaces documentation for specific actions under the MBeanPermission class.

- Use the following permission for WebSphere Application Server administrative application programming interfaces (APIs):

```
permission com.ibm.websphere.security.WebSphereRuntimePermission "AdminPermission";
```

Java Management Extensions MBean multiprocess model request flow for WebSphere Application Server for z/OS

Using the Java Management Extensions (JMX) dynamic proxy capability, applications that depend on JMX operations can exhibit consistent behavior whether the server architecture uses a single process model or a multiprocess model.

All the MBeans that the WebSphere Application Server runtime provides are capable of running under the single process model employed by WebSphere Application Server on distributed platforms, or the multiprocess model employed by WebSphere Application Server for z/OS. User MBean providers might need to modify their MBeans so that they work on both the WebSphere Application Server distributed platforms and the WebSphere Application Server for z/OS. For more information, view the application programming interfaces documentation.

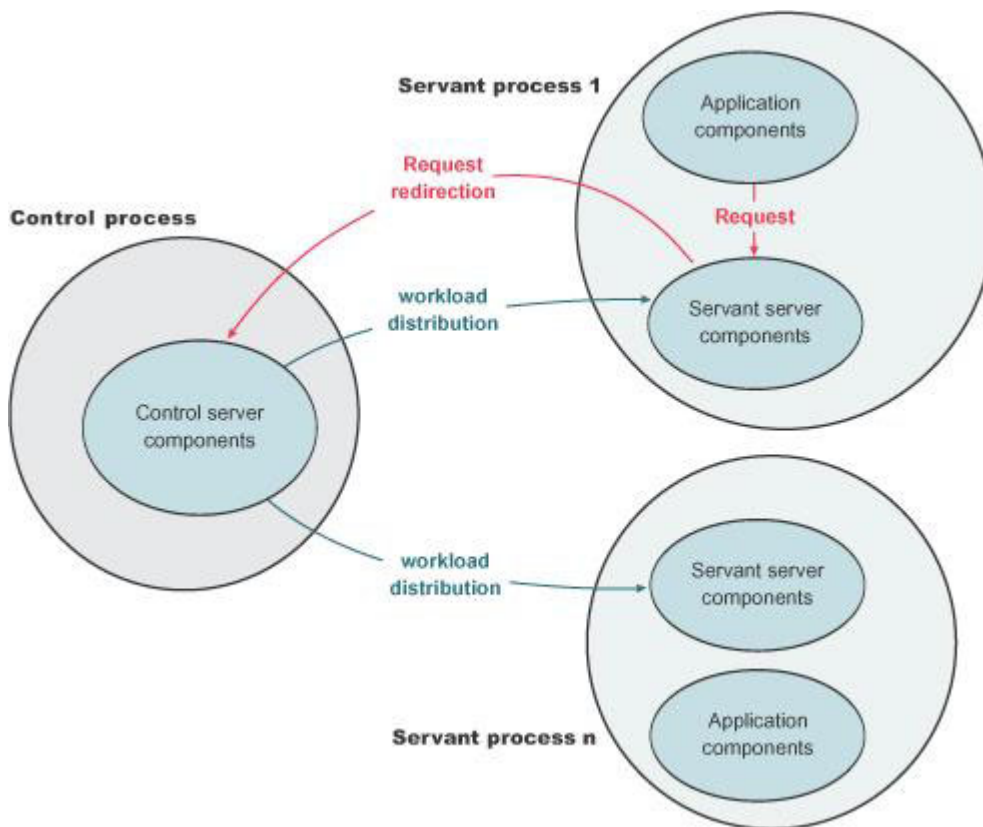
The simplified dynamic proxy model presented here discusses the two general request flows that exist in the multiprocess model. Operation requests on an MBean can be initiated from one of two places:

- Within an application component running under the same servant process as the MBean.

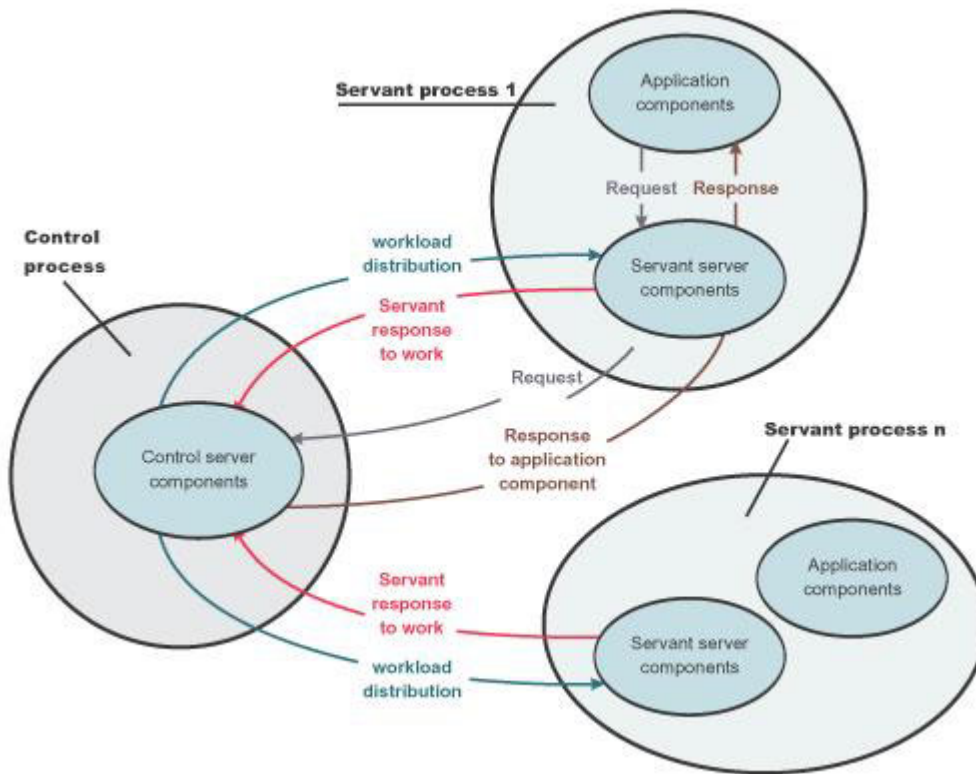
- Outside the server through one of the JMX connectors (SOAP, Remote Method Invocation (RMI), HTTP, and so on).

Requests that come from an application component generally follow this flow. The application component sends its request to the servant components. The servant components redirect the request to the control process where the dynamic proxy for the MBean runs. WebSphere Application Server automatically generates the MBean dynamic proxy during runtime. If the request calls a method that the MBean provider defined with a unicast option, the dynamic proxy in the control process randomly dispatches the work, with MVS workload management (WLM), to one servant process. If the request calls a method that the MBean provider defined with a multicast option, the dynamic proxy in the control process, in conjunction with WLM, distributes the work to all the servant processes that have activated the MBean.

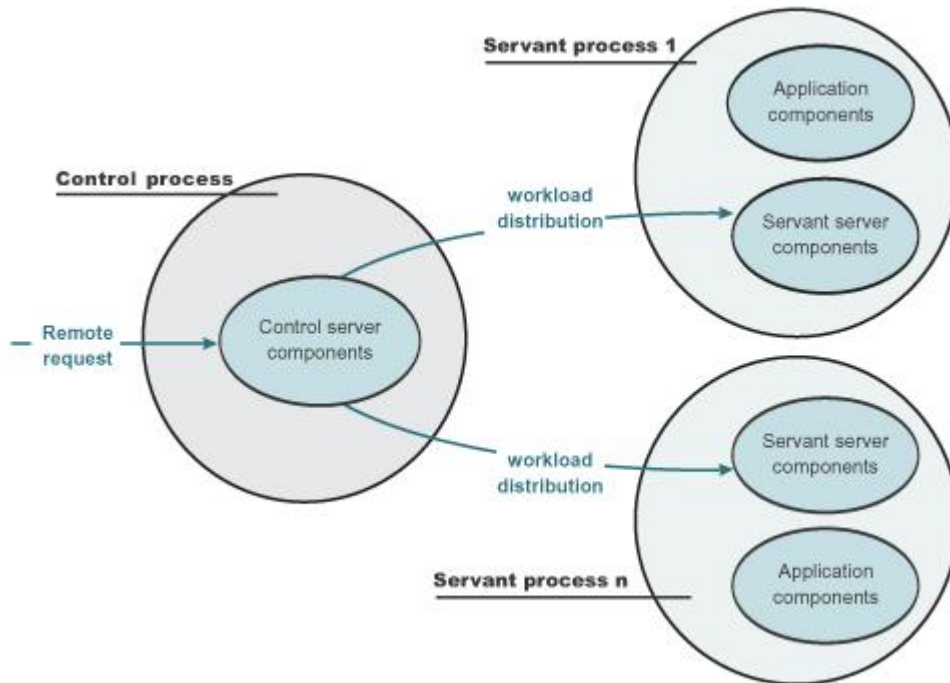
MBeans associated with application components are automatically registered with the dynamic proxy in the control process as part of the MBean activation process when WLM starts the servant process. MBeans that are activated in a single servant process, such as an MBean that is created and registered dynamically by an application component, cannot participate in multicast operations. Refer to the topic Extending the WebSphere Application Server administrative system with custom MBeans for information on how to register an MBean.



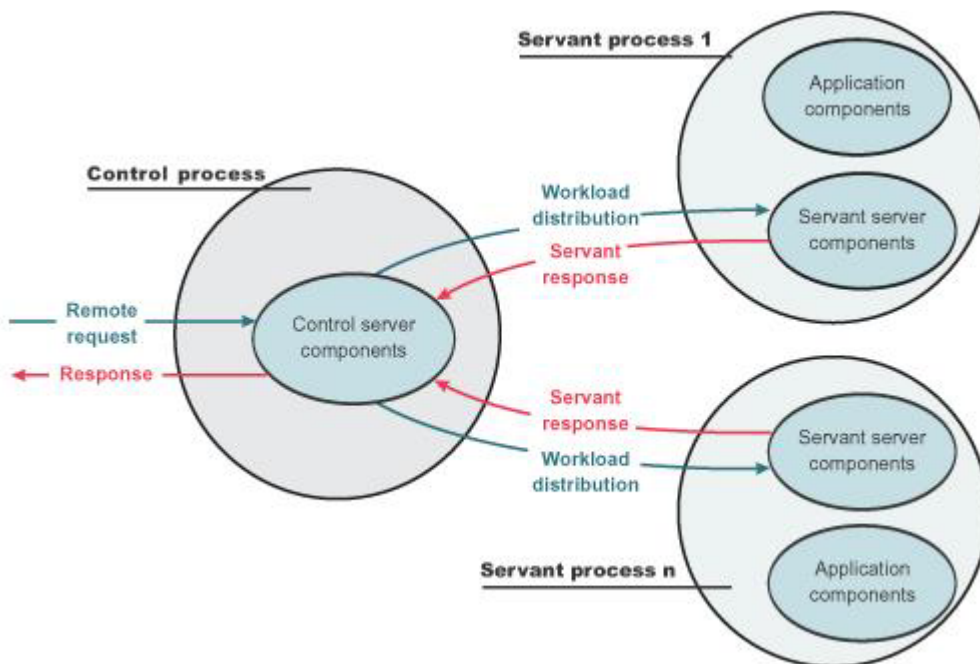
Each servant process that completes the work optionally sends a response back to the control process. If the MBean was defined with a unicast option, and the return type is anything but void, the control process returns the response to the servant that made the request. The servant server components then return the response to the application component. If the MBean was defined with a multicast option, the MBean inside each servant process runs separately and finishes processing the request at different times. After all the requests are processed, you might need a result aggregation and an event aggregation to properly return a result to the application component.



Requests that come from a remote location outside the server generally follow this flow. The remote location sends its request to the control process where the dynamic proxy for the MBean runs. WebSphere Application Server automatically generates the MBean dynamic proxy during run time. The next part of the request process behaves the same as a request that originates from an application component. If the request calls a method that the MBean provider defined with a unicast option, the dynamic proxy in the control process randomly dispatches the work, with MVS workload management (WLM), to one of the servant process in the list of servant processes that have activated the MBean. If the request calls a method that the MBean provider defined with a multicast option, the dynamic proxy in the control process, in conjunction with WLM, distributes the work to all the servant process that have activated the MBean.



The flow of the response back to the remote location is similar to the response back to the application component. Each servant process that completes the work optionally sends a response back to the control process. If the MBean was defined with a unicast option, and the return type is anything but void, the control process returns the response to the remote location. If the MBean was defined with a multithreaded option, the MBean that runs inside each servant process runs separately and finishes processing the request at different times. After all the requests are processed, you might need a result aggregation and an event aggregation to properly return a result to the remote location.



Java Management Extensions dynamic proxy concepts:

A Java Management Extensions (JMX) dynamic proxy coordinates MBean requests among multiprocess servers. This section discusses the main terms associated with a JMX dynamic proxy.

Control process

Receives requests and distributes them to servant processes so that the application server can do work for the requests.

Servant process

Receives work from the control process and carries out the work.

Unicall option versus the multicall option

Use the unicall option on the proxyInvocationType method when a request invokes an arbitrary servant process or servant processes. Use the multicall option on the proxyInvocationType method when a request goes to multiple servant processes and the servant processes return different results.

The following example shows an MBean descriptor that was developed for a single process model (before) and modified for a multiprocess model (after).

Before

```
<?xml version="1.0" encoding="UTF-8"?>
<!DOCTYPE MBean SYSTEM "MbeanDescriptor.dtd">
<MBean type="SampleStateMBean"
  version="6.0"
  description="Sample State MBean for the documentation example.">

  <attribute description="The name of the MBean."
    getMethod="getMBeanName" name="mbeanName" type="java.lang.String"/>
  <attribute description="The state of the MBean." name="state"
    getMethod="getState" setMethod="setState" type="java.lang.String"/>
  <operation
    description="Initialize the State MBean."
    impact="ACTION" name="initializeState" role="operation"
    targetObjectType="objectReference" type="void">
    <signature>
      <parameter description="The name of the MBean."
        name="mbeanName" type="java.lang.String"/>
      <parameter description="The initial state of the MBean."
        name="mbeanName" type="java.lang.String"/>
    </signature>
  </operation>
</MBean>
```

After

```
<?xml version="1.0" encoding="UTF-8"?>
<!DOCTYPE MBean SYSTEM "MbeanDescriptor.dtd">
<MBean type="SampleStateMBean"
  version="6.0"
  platform="dynamicproxy"
  description="Sample State MBean for the documentation example.">

  <attribute description="The name of the MBean."
    getMethod="getMBeanName" name="mbeanName" type="java.lang.String"/>
  <attribute description="The state of the MBean." name="state"
    getMethod="getState" setMethod="setState" type="java.lang.String"/>
  proxyInvokeType="unicall" proxySetterInvokeType="multicall"/>
  <operation
    description="Initialize the State MBean."
    impact="ACTION" name="initializeState" role="operation"
    targetObjectType="objectReference" type="void" proxyInvokeType="multicall">
    <signature>
      <parameter description="The name of the MBean."
        name="mbeanName" type="java.lang.String"/>
      <parameter description="The initial state of the MBean."
        name="mbeanName" type="java.lang.String"/>
    </signature>
  </operation>
```

```

        name="mbeanName" type="java.lang.String"/>
    </signature>
</operation>
</MBean>

```

Make the user MBean run in dynamic proxy mode by specifying `dynamicproxy` on the `platform` attribute. If no `platform` attribute exists on the MBean descriptor, the user MBean deployed on WebSphere Application Server for z/OS automatically uses the dynamic proxy mode.

Update the attribute XML tag or the operation XML tag, as shown in the After example, to specify the `unicall` behavior or the `multicall` behavior in the multiprocess environment. If no `proxyInvokeType` option or `proxySetterInvokeType` option exists, the behavior defaults to one of the following values:

```

proxyInvokeType=unicall (for the getMethod)
proxyInvokeType=multicall (for the setMethod)

```

In the After example, the `getMBeanName` method and the `getState` method run with `unicall` behavior. The `setState` method and the `initializeState` method run with `multicall` behavior.

Single process model

The single process application server has one server run time. The MBean generally acts on one instance of each major run time component: one Enterprise JavaBeans (EJB) container, one web container, one Java 2 Platform, Enterprise Edition (J2EE) connection manager, and so on. This model assumes that each MBean invocation on the server runs in the same process and the same Java Virtual Machine (JVM).

Multiprocess model

The multiprocess model asserts that a single server instance is a federation of Java virtual machines (JVMs), that run in a separate operating process. The control process is responsible for such server functions as communication endpoints, authorization, resource recovery, and workload management. All other JVMs are worker JVMs, in which application requests run. These JVMs take direction from, and interact only with the control process.

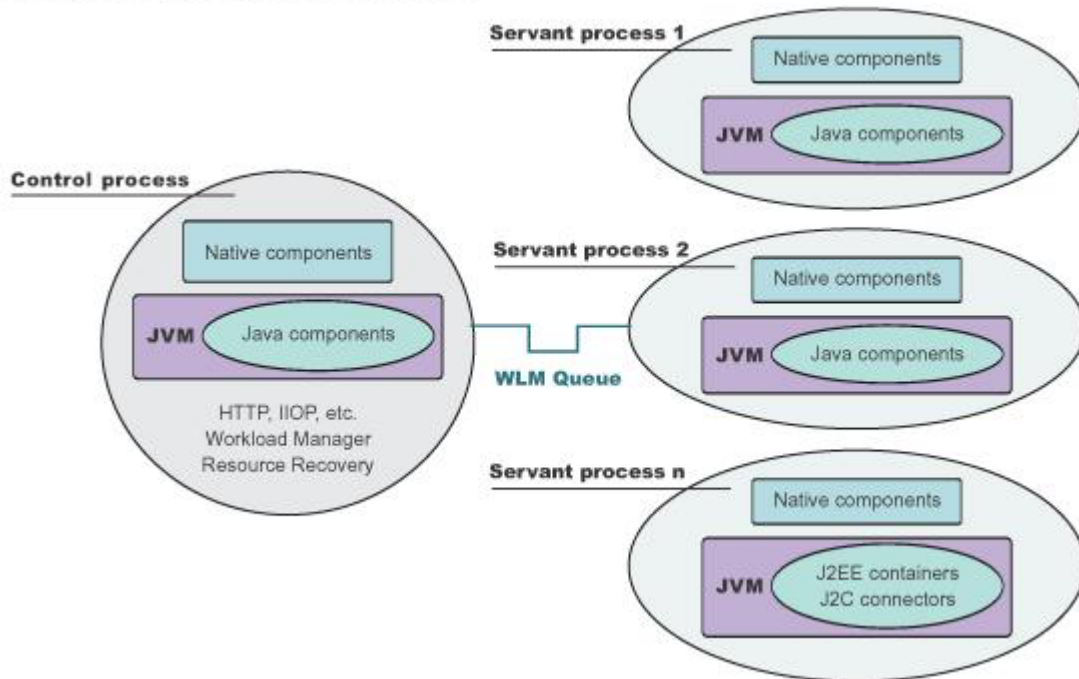
All inbound and outbound requests go through the control process. Client requests arrive at the control process. The control process, with assistance from the MVS workload manager (WLM), dispatches the work to the servant processes.

The number of servant processes is managed by WLM and varies based on demand. The demand is measured against installation-specific performance goals, expressed as WLM policy. Each servant process is identical and hosts the necessary application server components to enable the J2EE application programming model. The servant processes rely on the control process for numerous services, such as communication, security, and transaction control.

The multiprocess model imposes additional demands on the Java Management Extension (JMX) infrastructure over the single process model. Administrative requests to a multiprocess server often require coordination among the processes that comprise the application server. The JMX infrastructure

includes additional facilities to enable this coordination.

WebSphere Application Server for z/OS



State object support for dynamic proxy MBean

com.ibm.websphere.management.dynamicproxy.StateObject class: The MBean provider extends the `StateObject` abstract class. Specify the subclass of the `StateObject` class so that the JMX run time can instantiate it before the dynamic proxy MBean completes its initialization. The JMX run time attaches `StateObject` class to the dynamic proxy `InvocationHandler` interface to keep track of the current state of the dynamic proxy before and after the MBean method runs. The JMX run time also attaches the `StateObject` class to the `Result Aggregation` interface class as well as the `Event Handler` interface class to support appropriate aggregation application.

Result aggregation handler support interface

com.ibm.websphere.management.dynamicproxy.AggregationHandler class: The result aggregation handler support interface defines the method that an MBean provider uses to handle result aggregation in a dynamic proxy-enabled WebSphere Application Server for z/OS MBean. Specify the `aggregationHandlerClass` attribute on the `MBeanDescriptor` MBean XML tag. Implement the interface for MBean methods that use the `multicall proxyInvokeType` option and that return a value. The interface determines the method for which this aggregation is processed. It then properly aggregates all servant MBean results that the servant processes pass back to the control process, and then compiles a single result to return to the caller.

Event aggregation handler support interface

com.ibm.websphere.management.dynamicproxy.EventHandler interface class: The event aggregation handler support interface defines the method that an MBean provider uses to handle event aggregation in a dynamic proxy-enabled WebSphere Application Server for z/OS MBean. Specify the `eventHandlerClass` attribute on the `MBeanDescriptor` MBean XML tag. The interface handles all incoming servant MBean events and aggregates them to filter out duplicate events from multiple servant MBeans. It sends one event back to the listener of the dynamic proxy MBean. The interface adjusts the current dynamic proxy MBean state according to the MBean provider requirements.

Invocation handler support interface

com.ibm.websphere.management.dynamicproxy.InvocationHandler class: The invocation handler support interface defines the `preInvoke` and `postInvoke` methods that a WebSphere Application Server for z/OS dynamic proxy MBean implements when it requires state management information. The

MBean uses the information to coordinate with the servant MBeans in cases where the multicall invocation type is required. Specify the invocationHandlerClass attribute on the MBeanDescriptor MBean XML tag. Use the interface for dynamic proxy MBeans that require state management before and after invoking a method that changes its state.

User MBean

The user MBean resides in the servant process and handles requests through its dynamically created proxy MBean, which runs inside the control process. An MBean provider can package handlers with the user MBean so that the provider hooks in his own specialized processing for the following situations:

- Result aggregation
- Event aggregation
- Invocation handling
- State management of objects

Example: The SampleStateMBean MBean:

Use this example to guide you in developing user MBeans that work for the WebSphere Application Server on both the distributed platforms and the z/OS platform. The example uses all the special handlers to show its dynamic proxy MBean responsibilities and capabilities. The SampleStateMBean example keeps track of its state and generates state change events when it invokes setter methods.

MBeanDescriptor

```
<?xml version="1.0" encoding="UTF-8"?>
<!DOCTYPE MBean SYSTEM "MbeanDescriptor.dtd">
<MBean type="SampleStateMBean"
  aggregationHandlerClass="com.ibm.ws390.sample.SampleStateAggregationHandler"
  eventHandlerClass="com.ibm.ws390.sample.SampleStateEventHandler"
  invocationHandlerClass="com.ibm.ws390.sample.SampleStateInvocationHandler"
  stateObjectClass="com.ibm.ws390.sample.SampleState"
  version="6.0"
  platform="dynamicproxy"
  description="Sample State MBean for the documentation example.">

  <attribute description="The name of the MBean."
    getMethod="getMBeanName" name="mbeanName" type="java.lang.String"
    proxyInvokeType="unicall"/>
  <attribute description="The state of the MBean." name="state"
    getMethod="getState" setMethod="setState" type="java.lang.String"
    proxyInvokeType="multicall" proxySetterInvokeType="multicall"/>
  <operation
    description="Initialize the State MBean."
    impact="ACTION" name="initializeState" role="operation"
    targetObjectType="objectReference" type="void" proxyInvokeType="multicall">
    <signature>
      <parameter description="The name of the MBean."
        name="mbeanName" type="java.lang.String"/>
      <parameter description="The initial state of the MBean."
        name="mbeanName" type="java.lang.String"/>
    </signature>
  </operation>

  <notification name="j2ee.state.starting" severity="6" log="false"
    description="This sample state MBean is in starting state.">
    <notificationType>j2ee.state.starting</notificationType>
  </notification>
  <notification name="j2ee.state.running" severity="6" log="false"
    description="This sample state MBean is in running state.">
    <notificationType>j2ee.state.running</notificationType>
  </notification>
  <notification name="j2ee.state.stopping" severity="6" log="false"
    description="This sample state MBean is in stopping state.">
```

```

        <notificationType>j2ee.state.stopping</notificationType>
</notification>
<notification name="j2ee.state.stopped" severity="6" log="false"
        description="This sample state MBean is in stopped state.">
        <notificationType>j2ee.state.stopped</notificationType>
</notification>
</MBean>

```

SampleState implementation

```

package com.ibm.ws390.sample;

import com.ibm.ejs.ras.Tr;
import com.ibm.ejs.ras.TraceComponent;
import java.io.Serializable;
import com.ibm.websphere.management.dynamicproxy.StateObject;

public class SampleState extends StateObject {

    private static TraceComponent tc =
        Tr.register(SampleState.class,"SampleState",null);

    // Package protected STATE constants.
    static final String STATE_STARTING = "j2ee.state.starting";
    static final String STATE_RUNNING = "j2ee.state.running";
    static final String STATE_STOPPING = "j2ee.state.stopping";
    static final String STATE_STOPPED = "j2ee.state.stopped";

    // Dynamicproxy State is initialized with STOPPED state.
    private String state = STATE_STOPPED;

    public SampleState() {
        if (tc.isEntryEnabled()) Tr.entry(tc,"<init>");

        // State is initialized during "state" initialization above,
        // but can also be initialized here in the constructor as well.
        /*
        state = "WebSphere Application Server for z/OS ready for e-business";
        */

        if (tc.isEntryEnabled()) Tr.exit(tc,"<init>");
    }

    public synchronized String getState() {
        if (tc.isEntryEnabled()) Tr.entry(tc,"getState");
        if (tc.isEntryEnabled()) Tr.exit(tc,"getState",state);
        return state;
    }

    public synchronized void setState(String state) {
        if (tc.isEntryEnabled()) Tr.entry(tc,"setState",state);
        this.state = state;
        if (tc.isEntryEnabled()) Tr.exit(tc,"setState");
    }

    public synchronized String getStateObjectInfo() {
        return state;
    }
}

```

SampleStateAggregationHandler implementation

```

package com.ibm.ws390.sample;

import com.ibm.websphere.management.dynamicproxy.AggregationHandler;
import com.ibm.websphere.management.dynamicproxy.StateObject;

import com.ibm.ejs.ras.Tr;

```

```

import com.ibm.ejs.ras.TraceComponent;

public class SampleStateAggregationHandler implements AggregationHandler {

    private static TraceComponent tc =
    Tr.register(SampleStateAggregationHandler.class,"SampleState",null);

    /**
     * Return an aggregated result from a multicall Mbean operation which
     * compiles through all servant MBeans' results and returns a respective
     * single return value for an invoked method.
     *
     * @param methodName      MBean method name
     * @param params          MBean method parameters
     * @param signatures      MBean method signatures
     * @param servantMBeanResults Result of each servant MBean instances
     *                        invoked by the dynamicproxy multicast
     *                        invocation.
     *
     * Note: this value can be "null" OR can be
     * an array of "null"s in case return value
     * of the method is "void." Implementation
     * of this method MUST handle this case to
     * avoid a <code>NullPointerException</code>.
     *
     * @param stateObject
     * MBean provider provided <code>StateObject</code> used by
     * dynamicproxy MBean in CR to manage its state. Note: this object
     * MAY BE null if "stateObjectClass" was not specified OR internal
     * error occurred during initialization of this dynamicproxy MBean.
     * Implmentation MUST properly handle "null" input.
     *
     * @return aggregated result as defined by MBean xml for specified
     * MBean operation.
     */
    public Object aggregateResults(String methodName,
                                  Object[] params,
                                  String[] signatures,
                                  Object[] servantMBeanResults,
                                  StateObject stateObject) {

        if (tc.isEntryEnabled()) Tr.entry(tc,"aggregateResults",methodName);

        // As you can see from the MBeanDescriptor of SampleStateMBean,
        // it declares the following four methods:
        // 1. String getMBeanName()          [proxyInvokeType == unicast]
        // 2. String getState()              [proxyInvokeType == multical]
        // 3. void setState(String)          [proxyInvokeType == multical]
        // 4. void initializeState()         [proxyInvokeType == multical]
        //
        // Looking at the above methods, only method that requires aggregation
        // is #2 getState method which is a multical MBean operation AND
        // it returns a value that can be aggregated.
        //
        // In this example, we simply take each servants' getState MBean
        // request result and concatenate them into one long String that
        // displays each servants' state.
        if (methodName.equals("getState")) {
            StringBuffer stateBuf = new StringBuffer();

            for (int i=0; i<servantMBeanResults.length; i++) {
                stateBuf.append("SERVANT #" + i + " state ==|" +
                               servantMBeanResults[i] + "|== ");
            }
            return stateBuf.toString();
        }
        // If we also had an example method which returns say an int,
        // getNumberOfMBeans(), it can take the similar approach
        // and to add each servants' getNumberOfMBeans() result together here.
    }
}

```

```

        /* example added for non-existent method: int getNumberOfMBeans()
        else if (methodName.equals("getNumberOfMBeans")) {
            int aggregatedResult = 0;

            for (int i=0; i<servantMBeanResults.length; i++) {
                aggregatedResult += (int) servantMBeanResults[i];
            }
            return aggregatedResult;
        }
        */

        return methodName + " is NOT handled by " + getClass().getName() + "!";
    }
}

```

Administrative security

Access to the Java Management Extension (JMX) administrative subsystem requires role-based access control when administrative security is enabled.

. A client, which can be a user or an administrative client program, can access an MBean method only if at least one of the required roles is granted to the client. WebSphere Application Server uses the declarative security approach to specify the security policy on the JMX MBean. This approach has the advantage of not requiring MBean developers to add security code. Moreover, WebSphere Application Server provides a default security policy for an MBean so in most case MBean developers do not need to specify a security policy at all. With WebSphere Application Server, you can define explicit security policy for your MBeans if the default security policy does not meet your specific security requirements.

Default MBean security policy

This topic discusses the default managed bean (MBean) security policy. In most cases, MBean developers do not need to specify a security policy.

Three types of MBeans exist for the default MBean security policy:

- A configuration type MBean
- A runtime type MBean
- A deployer type MBean

An optional attribute in the MBean descriptor XML file defines the type of MBean.

The ConfigRepository MBean is an example of one of a few configuration types. In the configRepository.xml descriptor file, the configureMBean = "true" attribute indicates that the MBean is a configuration type.

```

<MBean type="ConfigRepository"
  version="5.0"
  platform="common"
  description="Management interface for the configuration repository."
  configureMBean="true">

```

The EJBModule MBean is an example of deployer type MBeans. In the EJBModule.xml descriptor file, the deployerMBean="true" attribute indicates that the MBean is a deployer type.

```


<MBean type="EJBModule" j2eeType="EJBModule"
  version="5.0"
  platform="dynamicproxy"
  resourceIdentifierKey="Application"
  resourceType="Application"
  deployerMBean="true"
  description="Management interface for the EJBModule component.">

```

WebSphere Application Server extended role-based access control supports role inheritance. Five administrative roles of administrator, configurator, operator, deployer, and monitor exist. The monitor role is the least privileged administrative role. Users that are granted the monitor role can view the WebSphere

Application Server configuration and the runtime status, but cannot make any changes. The other administrative roles each have their own unique set of privileges as well as the same privileges as the monitor role.

The configurator role has permission to modify WebSphere Application Server configuration data. The operator role has permission to change the runtime state, such as the start and stop of administrative resources. A configurator role cannot change the runtime status and conversely an operator role cannot change the WebSphere Application Server configuration. The administrator role includes configurator and operator role, but has more permissions than the union of configurator role and operator role. The administrator role can additionally change the administrative security configuration. A simple picture shows the administrative role inheritance relationship. The deployer role is a combination of the configurator and operator roles for application management. The deployer role has both configurator and operator permission for applications. A diagram shows the administrative role inheritance relationship.



Each MBean method or operation is assigned an impact attribute with a value of either INFO or ACTION. Here are some examples:

- A get method has an impact value of INFO and a write method has an impact value of ACTION.
- In the ConfigRepository MBean, the extract method does not change the configuration data and has an impact value of INFO, while the modify method has an impact value of ACTION.
- In the Java virtual machine (JVM) MBean, which is an operator type of MBean, the ggetCurrentTimeInMillis method has an impact value of ACTION.

A configuration MBean method that has an impact value of INFO requires the monitor role. A configuration MBean method that has an impact value of ACTION requires the configurator role. A deployer MBean method that has an impact value of INFO requires the monitor role. A deployer MBean method that has an impact value of ACTION requires the deployer role. Because all administrative roles are monitor roles, any administrative role can access configuration MBean methods and deployer MBean methods that have an impact value of INFO. The administrator role is a configurator role and has access to the configuration MBean methods that have an impact value of ACTION.

The default security policy for the configuration MBean is summarized in the following table:

Table 8. Configuration MBean method impact values and security roles. An X indicates that the MBean method requires the role by default.

Method impact	Monitor role	Operator role	Configurator role	Deployer role	Administrator role
INFO	X	X	X	X	X
ACTION			X		X

The default security policy for the operation MBean is summarized in the following table:

Table 9. Operation MBean method impact values and security roles. An X indicates that the MBean method requires the role by default.

Method impact	Monitor role	Operator role	Configurator role	Deployer role	Administrator role
INFO	X	X	X	X	X
ACTION		X			X

The default security policy for the deployer MBean is summarized in the following table:

Table 10. Deployer MBean method impact values and security roles. An X indicates that the MBean method requires the role by default.

Method impact	Monitor role	Operator role	Configurator role	Deployer role	Administrator role
INFO	X	X	X	X	X
ACTION		X		X	X

If an MBean has both the `configureMBean` attribute and the `deployerMBean` attribute set to `true`, the required role for all actions is either configurator or monitor. No such MBean is presently defined in the system.

Defining an explicit MBean security policy

You can explicitly define an MBean security policy for a particular MBean. Use this example to define an MBean security policy.

Before you begin

This task assumes a basic familiarity with MBean programming. For information on MBean programming, see MBean Java application programming interface (API) documentation.

About this task

Perform the following tasks to define an explicit security policy.

Procedure

1. Assume that you have an MBean defined by the MBean `sample.xml` descriptor file.
2. Specify the explicit security policy for that MBean in the `sampleSecurity.xml` file. The naming convention is that you must append "Security" to the MBean descriptor file name as the name of the MBean security descriptor file.

3. Place the security policy descriptor file at the same directory where the MBean security descriptor file is so that the MBean loader can find it. This directory is the typical location for the security policy descriptor file. If no MBean security descriptor file is present, the default MBean security policy is used.
4. Specify the MBean name of `sample` in the resource element `resource-name` field of the `sampleSecurity.xml` file so that the MBean policy loader can associate the MBean security policy with the MBean. The MBean security descriptor definition is very similar to the security policy that is defined by the Java 2 Platform, Enterprise Edition (J2EE) deployment descriptor.

Results

You now have an explicitly defined MBean security policy that you can run with an MBean.

Example

The following example describes the MBean security descriptor file format for the `sampleSecurity.xml` file.

Line 2 specifies that an MBean security descriptor schema is defined by the `RolePermissionDescriptor.dtd` file, which is a document type definition (DTD) in WebSphere Application Server.

As shown on line 3, each MBean descriptor file contains a single `role-permission` element. The administrative security role hierarchy is defined in the `security-role` elements between line 9 and line 37. The administrative security role has an inheritance relationship.

As defined on line 14 through 21, the operator security role implies the monitor security role, which means that a user with the operator role has all the permissions of the monitor role. As defined between line 30 and line 38, an administrator security role implies both the configurator and operator security role. Every MBean security descriptor file typically has the same role relationship definition so that you can cut and paste this section to your MBean security descriptor file.

One or more `method-permission` elements are defined after the `security-role` element. Each `method-permission` element defines the required roles for one or more methods. Specify method parameters to avoid method name collision in case multiple methods have the same name.

```

1. <?xml version="1.0" encoding="UTF-8"?>
2. <!DOCTYPE role-permission SYSTEM "RolePermissionDescriptor.dtd" >
3. <role-permission>
4.   <resource>
5.     <resource-name>sample</resource-name>
6.     <class-name>com.ibm.ws.security.descriptor.sample</class-name>
7.     <description>This is a sample for testing role permission descriptor.</description>
8.   </resource>
9.   <security-role>
10.    <role>
11.      <role-name>monitor</role-name>
12.    </role>
13.  </security-role>
14.  <security-role>
15.    <role>
16.      <role-name>operator</role-name>
17.      <imply>
18.        <role-name>monitor</role-name>
19.      </imply>
20.    </role>
21.  </security-role>
22.  <security-role>
23.    <role>
24.      <role-name>configurator</role-name>
25.      <imply>
26.        <role-name>monitor</role-name>
27.      </imply>

```



```

28.     /role>
29. </security-role>
30. <security-role>
31.     <role>
32.         <role-name>administrator</role-name>
33.         <imply>
34.             <role-name>operator</role-name>
35.             <role-name>configurator</role-name>
36.         </imply>
37.     </role>
38. </security-role>
39. <method-permission>
40.     <description>Sample method permission table</description>
41.     <role-name>operator</role-name>
42.     <method>
43.         <description>Sample operation</description>
44.         <resource-name>sample</resource-name>
45.         <method-name>stop</method-name>
46.     </method>
47. </method-permission>
48. <method-permission>
49.     <description>Sample method permission table</description>
50.     <role-name>operator</role-name>
51.     <method>
52.         <description>Sample operation</description>
53.         <resource-name>sample</resource-name>
54.         <method-name>start</method-name>
55.         <method-params>
56.             <method-param>java.lang.String</method-param>
57.             <method-param>java.lang.String</method-param>
58.         </method-params>
59.     </method>
60. </method-permission>
61. <method-permission>
62.     <description>Sample method permission table</description>
63.     <role-name>operator</role-name>
64.     <method>
65.         <description>Sample operation</description>
66.         <resource-name>sample</resource-name>
67.         <method-name>monitor</method-name>
68.         <method-params>
69.         </method-params>
70.     </method>
71. </method-permission>
72. <method-permission>
73.     <description>Sample method permission table</description>
74.     <role-name>configurator</role-name>
75.     <method>
76.         <description>Sample operation</description>
77.         <resource-name>sample</resource-name>
78.         <method-name>setValue</method-name>
79.         <method-params>
80.             <method-param>java.lang.Boolean</method-param>
81.         </method-params>
82.     </method>
83. </method-permission>
84. <method-permission>
85.     <description>Sample method permission table</description>
86.     <role-name>monitor</role-name>
87.     <method>
88.         <description>Sample operation</description>
89.         <resource-name>sample</resource-name>
90.         <method-name>getValue</method-name>
91.     </method>
92. </method-permission>
93. </role-permission>

```

Specifying fine-grained MBean security in the MBean descriptor

To implement fine-grained administrative security, your code must identify the resource instance that the managed bean (MBean) represents and assign the user the required role for that instance of the resource. This topic discusses what to do to identify the resource and assign the required role. This topic also discusses how to make an MBean method run under a different user identity so that the method can access other resource instances. Lastly, this topic discusses how to check if an MBean method has access to a resource instance by using programmatic interfaces.

Before you begin

This task assumes a basic familiarity with MBean programming. For information on MBean programming, see MBean Java application programming interface (API) documentation.

About this task

Perform the following task to ensure that an MBean or MBean method is protected. Identify the resource instance that the MBean or MBean method represents and assign required roles to access the MBean. Perform this task during the development of the MBean.

Procedure

1. Determine the resource instance that the MBean represents and the required roles to invoke the MBean methods.

Every MBean method has a default MBean security policy. When the MBean method uses the default security policy, the resource instance that the MBean represents is assumed to be the server in which the MBean runs. If an MBean or MBean method represents a resource instance other than the server on which it runs, perform the following steps:

- a. Identify the resource instance that the MBean represents.
 - If an MBean, such as the Server MBean, accesses and modifies the server in which the MBean runs, do not specify a security policy to verify that the user invoking the MBean is granted access to the server because the default security policy is in force. In most cases, you use an MBean to access and modify the server.
 - If an MBean that runs inside a server can access and modify resources that do not directly belong to the server, check if the user invoking the MBean is granted access to the instance of the resource before allowing the MBean method to run.

In most cases, identify the resource instance by identifying the key-value pair in the object name of the MBean that represents the resource instance. The `resourceIdentifierKey` attribute defines the key.

For example, you can use the `EJBModule` MBean to access an Enterprise JavaBeans (EJB) module within an application that runs inside the server. In this case, the object name of the `EJBModule` MBean contains a key-value pair. The key is `Application`. The value represents the resource instance that the `EJBModule` MBean tries to access. The user that invokes this MBean method is verified to make sure that access is granted to this instance of the application before allowing the MBean method to run.

The following example shows how to describe the fine-grained administrative security for the `EJBModule` type of MBean in the MBean descriptor:

```
<?xml version="1.0" encoding="UTF-8"?>
<!DOCTYPE MBean SYSTEM "MbeanDescriptor.dtd">
<MBean type="EJBModule" j2eeType="EJBModule"
  version="5.0"
  platform="dynamicproxy"
  resourceIdentifierKey="Application"
  resourceType="Application"
  deployerMBean="true"
  description="Management interface for the EJBModule component.">
```

- If you can determine the resource that the MBean accesses before the MBean is invoked, but you cannot use the MBean object name to determine the resource instance that the MBean accesses, use parameters that are passed to the MBean instead.

Identify the MBean method parameter name with a parameter value that represents the resource instance. Mark the corresponding parameter metadata in the MBean descriptor as the resource identifier. To mark a parameter as the resource identifier, add the `resourceType` attribute. The attribute specifies the type resource that the parameter value contains. When the `resourceType` attribute is present for any MBean method parameter, the parameter value determines the resource instance that the MBean method represents.

For example, one instance of the `ApplicationManager` MBean runs in each server. The same MBean can be used to start and stop all the applications in the server. The start and stop methods of this MBean each take the application name as a parameter. They use the parameter to determine the instance of the application that this MBean method tries to access.

The following example shows how to describe the fine-grained administrative security for this type of MBean in the MBean descriptor:

```
<operation
  description="Start Application"
  impact="ACTION" name="startApplication" role="operation"
  targetObjectType="objectReference" type="void" proxyInvokeType="spray">
  <signature>
    <parameter description="Application Name" resourceType="Application"
      name="applicationName" type="java.lang.String"/>
  </signature>
</operation>
```

- If the resource that an MBean accesses cannot be determined until the MBean is invoked, check if the user invoking the MBean is granted access to the instance of the resource by using application programming interfaces (APIs).

Mark the MBean or MBean method as excluded from access checking in the MBean descriptor by using the `excludeAccessCheck` attribute. When an MBean is marked as excluded from access checking, all its methods are also excluded from access checking.

For example, the `ConfigService` MBean that runs in the deployment manager is used to configure all the resources within a cell. Exclude this MBean from access checking before invoking the MBean methods. Check that the `ConfigService` MBean is granted access to the configuration resource when the MBean attempts to access the resource.

The following example shows how to describe the fine-grained administrative security for the `ConfigServices` type of MBean in the MBean descriptor:

```
<?xml version="1.0" encoding="UTF-8"?>
<!DOCTYPE MBean SYSTEM "MbeanDescriptor.dtd">
<MBean
  version="5.0"
  platform="proxy"
  collaboratorClass="com.ibm.ws390.management.proxy.ConfigServiceManager"
  description="Config Service component provides service of
configuration related tasks on top of configuration repository service."
  type="ConfigService"
  excludeAccessCheck="true"
  configureMBean="true">
```

Some statements are split on multiple lines for printing purposes.

The following example shows how to invoke the MBean method logic to perform authorization checking programmatically:

```
// Get administration authorizer.
AdminAuthorizer aa = AdminAuthorizerFactory.getAdminAuthorizer();
// Set the role that is required for this operation.
String role = com.ibm.ws.security.util.Constants.CONFIG_ROLE;
// Set the resource name.
// cells/cellName is optional.
```

```

String resource = "/nodes/" + nodeName + "/servers/" + serverName;
// Check access
if ( aa != null && !aa.checkAccess(resource, role) )
    // Disallow access.
    else
        // Allow access.

```

- b. Assign required roles for the MBean and MBean methods.

The required roles are automatically assigned, based on the type of MBean and the impact of the MBean method, as described in the topic on the default MBean security policy.

2. Specify delegation mode.

In some cases, after performing the initial access check, the MBean method might need to run under a different user identity so that it can access other resource instances. For example the syncNode operation in the CellSync MBean grants the user the operator role to the instance of the node being synchronized. The syncNode operation tries to access resources under the cell scope. The user might not have access to open files under the cell directory. The MBean must run as System after the initial access check so that the operation completes without any access denied problems.

Set the runAs attribute to System to specify delegation mode for an MBean or MBean method. When you set the runAs attribute for an MBean, the value applies to all MBean methods for that MBean.

The following example shows how to describe fine-grained administrative security for the CellSync type of MBean in the MBean descriptor.

```

<?xml version="1.0" encoding="UTF-8"?>
<!DOCTYPE MBean SYSTEM "MbeanDescriptor.dtd">
<MBean type="CellSync"
  version="5.0.1"
  platform="common"
  runAs="System"
  description="Management interface for the configuration synchronization logic
  performed at the central deployment manager for the cell.">
  <operation
    description="Initiate a synchronization request for a given node" impact="ACTION"
    name="syncNode" role="operation" targetObjectType="objectReference" type="ja
va.lang.Boolean">
  <signature>
  <parameter resourceType="Node"
    description="The name of the node"
    name="nodeName" type="java.lang.String"/>
  </signature>
</operation>

```

Results

You have determined the type of resource a given MBean method is accessing and performed the necessary access check so that the product can allow access to the resource.

Administrative programs for multiple Java Platform, Enterprise Edition application servers

You can develop an administrative client to manage multiple vendor application servers through existing MBean support in the WebSphere Application Server.

Existence of MBeans for stopped components

best-practices: The WebSphere Application Server completely implements the Java Platform, Enterprise Edition (Java EE) Management specification. However, some differences in details between the Java EE specification and the WebSphere Application Server implementation are important for you to understand when you access WebSphere Application Server

components. These differences are important to you when you access application MBeans because you can use either the WebSphere Application Server programming model or the Java EE programming model.

In the WebSphere Application Server programming model, if an MBean exists, you can assume that it is running. If an MBean does not exist, you can assume that it is stopped. Transient states between the started state and the stopped state are the same as the stopped state, which means that no MBean exists.

In the Java EE programming model, the MBean always exists regardless of the state of the component.

You can determine the state of a component by querying the state attribute. However, the state attribute only exists for MBeans that are state manageable, meaning that they implement the `StateManageable` interface. State manageable MBeans have `start()`, `startRecursive()`, and `stop()` operations whether these MBeans are Java EE MBeans or WebSphere Application Server MBeans. Additionally, the WebSphere Application Server defines the stateful interface. The stateful interface means that the component has a state and emits the `Java EE.state.notifications` method, but that the component cannot directly manage the state. For example, a web module cannot stop itself. However, the application that contains the web module can stop it.

Not all MBeans that have a state are state-manageable. Servlets, Java EE modules and enterprise beans, for example, are all stateful, but are not state manageable. The Java EE server is not state-manageable because no `start()` operation is available on a server.

The `J2EEApplication` MBean is an example of a state manageable MBean. When the WebSphere Application Server starts, each application activates a `J2EEApplication` MBean for itself. A `J2EEApplication` MBean has a Java EE type of `J2EEApplication` (for example, `ObjectName **:*,j2eeType=J2EEApplication`). If the application starts, it also activates an `Application` MBean with a type of `Application` (for example, `**:*,type=Application`). When the application changes state, the `Application` MBean is activated or deactivated. However, the `J2EEApplication` MBean is always activated. You can retrieve the application state changes by getting the state attribute.

The `modules` attribute on the `J2EEApplication` component returns an array of object names, one for every module in the application. The Application Server activates an MBean for each of these modules only after the Application Server starts the application. The managed enterprise bean `isRegistered(ObjectName)` method returns `false` if the application, and therefore the module, is not running.

All of the attributes that are defined in the Java EE management specification return valid values when the managed object stops. Other attributes and operations, for example those that are specifically defined for the Application Server, use the `com.ibm.websphere.management.exception.ObjectNotRunningException` exception if they are accessed when the object is stopped.

If you install the application while the server runs, the application installs the `J2EEApplication` MBean when the installation completes. Conversely, when the application uninstalls the `J2EEApplication` MBean, the application deactivates the MBean.

Mapping type properties

You can determine which MBeans have a `j2eeType` property and a WebSphere Application Server type property in their `ObjectName` property sets by going to the additional application programming interfaces documentation. Select the MBean interfaces subtopic, and then specific MBeans in the list. Look for `type=` and `j2eeType=`. You can use the type property to query for any MBeans. MBeans derived from the Java EE specification have an additional `j2eeType` property as part of their `ObjectName` property sets. You can also use the `j2eeType` property to query for MBeans.

Optional WebSphere Application Server interfaces

The EventProvider, StateManageable, and StatisticsProvider interfaces are optional interfaces that the Java EE Management specification defines. Which of the interfaces the product implements varies from MBean to MBean. Go to the additional application programming interfaces documentation to see which interfaces the product implements for a particular MBean. Select the MBean interfaces subtopic, and then a specific MBean in the list. Find All Parent MBeans. The interfaces that are implemented for the MBean follow All Parent MBeans. For example, the J2EEDomain MBean does not implement any of the interfaces, while the JVM MBean implements the StatisticsProvider interface.

Deploying and managing a custom Java administrative client program with multiple Java Platform, Enterprise Edition application servers

This section describes how to connect to a Java Platform, Enterprise Edition (Java EE) server, and how to manage multiple vendor servers.

Before you begin

The product completely implements the Java EE Management specification, also known as JSR-77 (Java Specification Requests 77). However, some differences in details between the Java EE specification and the WebSphere Application Server implementation are important for you to understand when you develop a Java administrative client program to manage multiple vendor servers. For information, see the Java Platform, Enterprise Edition (Java EE) Management Specification and the MBean Java application programming interface (API) documentation.

About this task

When your administrative client program accesses WebSphere Application Servers exclusively, you can use the Java APIs and WebSphere Application Server-defined MBeans to manage them. If your program needs to access both WebSphere Application Servers and other Java EE servers, use the API defined in the Java EE Management specification.

Procedure

1. Connect to a Java EE server.

Connect to a server by looking up the Management enterprise bean from the Java Naming and Directory Interface (JNDI). The Management enterprise bean supplies a remote interface to the MBean server that runs in the application server. The Management enterprise bean works almost exactly like the WebSphere Application Server administrative client, except that it does not provide WebSphere Application Server specific functionality. The following example shows how to look up the Management enterprise bean.

```
import javax.management.j2ee.ManagementHome;
import javax.management.j2ee.Management;

Properties props = new Properties();

props.setProperty(Context.PROVIDER_URL, "iiop://myhost:2809");
Context ic = new InitialContext(props);
Object obj = ic.lookup("ejb/mgmt/MEJB");
ManagementHome mejbHome = (ManagementHome)
    PortableRemoteObject.narrow(obj, ManagementHome.class);
Management mejb = mejbHome.create();
```

The example gets an initial context to an application server by passing the host and port of the Remote Method Invocation (RMI) connector. You must explicitly code the RMI port, in this case 2809. The lookup method looks up the `ejb/mgmt/MEJB` path, which is the location of the Management enterprise bean home. The example then creates the `mejb` stateless session bean, which you use in the next step.

2. Manage multiple vendor application servers.

After you create the `mejb` stateless session bean, you can use it to manage your application servers. Components from the application servers appear as MBeans, which the specification defines. These MBeans all have the `j2eeType` property. This property is one of a set of types that the specification defines. All of these types have a set of exposed attributes.

Use the following example to guide you in managing multiple vendor application servers. The example uses the Java virtual machine (JVM) MBean to determine what the current heap size is for the application server.

```
ObjectName jvmQuery = new ObjectName("*:j2eeType=JVM,*");
Set s = mejb.queryNames(jvmQuery, null);
ObjectName jvmMBean = (ObjectName) s.iterator().next();
boolean hasStats = ((Boolean) mejb.getAttribute(jvmMBean,
    "statisticsProvider")).booleanValue();
if (hasStats) {
    JVMStats stats = (JVMStats) mejb.getAttribute(jvmMBean,
        "stats");
    String[] statisticNames = stats.getStatisticNames();
    if (Arrays.asList(statisticNames).contains("heapSize")) {
        System.out.println("Heap size: " + stats.getHeapSize());
    }
}
```

The `queryNames()` method first queries the JVM MBean. The `getAttribute` method gets the `statisticsProvider` attribute and determine if this MBean provides statistics. If the MBean does, the example accesses the `stats` attribute, and then invokes the `getHeapSize()` method to get the heap size.

Results

The strength of this example is that the example can run on any vendor application server. It demonstrates that an MBean can optionally implement defined interfaces, in this case the `StatisticsProvider` interface. If an MBean implements the `StatisticsProvider` interface, you can see if an application server supports a particular statistic, in this case the heap size. The specification defines the heap size, although this value is optional. If the application server supports the heap size, you can display the heap size for the JVM.

Java Management Extensions V1.0 to Java Management Extensions V1.2 migration

You might need to migrate custom MBeans that are supplied by products other than the Application Server from Version 5 to Version 6.0 and later for full compatibility.

Each Java virtual machine (JVM) in WebSphere Application Server includes an embedded implementation of Java Management Extensions (JMX). In Application Server, Version 5, the JVMs contain an implementation of the JMX 1.0 specification. In Application Server, Version 6.0 and later, the JVMs contain an implementation of the JMX 1.2 specification. The JMX 1.0 implementation used in Version 5 is the `TMX4J` package that IBM Tivoli products supply. The JMX 1.2 specification used in Version 6.0 and later is the open source `mx4j` package. The JMX implementation change across the releases does not affect the behavior of the JMX MBeans in the Application Server. No Application Server administrative application programming interfaces (APIs) are altered due to the change from the JMX V1.0 specification to the JMX V1.2 specification.

The JMX V1.2 specification is compatible with the earlier JMX V1.0 specification. However, you might need to migrate custom MBeans that are supplied by products other than the application server from Version 5 to Version 6.0 and later. The primary concern for these custom MBeans is related to the values that are used in key properties of the JMX `ObjectName` class for the MBean. The open source `mx4j` implementation more stringently enforces property validation according to the JMX 1.2 specification. Test the custom MBeans that you deployed in Version 5 in Version 6.0 and later, to ensure compatibility. Full details of the JMX V1.2 specification changes from the JMX V1.0 specification are available in the JMX 1.2 specification.

Java Management Extensions (JMX) interoperability

Starting with Version 6.0, WebSphere Application Server implements Java Management Extensions (JMX) Version 1.2.

Differences between Version 6.0.x, and Version 6.1 and later

The product supports communication with earlier levels of the server through the SOAP connector. The earlier levels of the server cannot be more than two releases earlier. Version 6 and later servers can communicate with each other through an RMI connector, with one restriction. The `javax.management.MBeanInfo` class and its subclasses do not interoperate between Version 6.1.02 and version 6.0.2. You will receive a `java.lang.IllegalArgumentException` exception when attempting to send any instance of these classes between a Version 6.1 and Version 6.0.2. This restriction affects clients that perform the following operations.

- A wsadmin scripting client that tries to perform the following operations, for example in Jacl,:

```
$Help attributes MBeanObjectName
$Help operations MBeanObjectName
$Help notifications MBeanObjectName
```

- Any Java client that tries to call:

```
com.ibm.websphere.management.AdminClient.getMBeanInfo(ObjectName name);
com.ibm.websphere.management.AdminService.getMBeanInfo(ObjectName name);
javax.management.MBeanServer.getMBeanInfo(ObjectName name);
```

A serialization format mismatch exists between the JMX implementation in Version 6.1 and later and Version 6.0.x releases. When a Version 6.0.x wsadmin script or a Version 6.0.x administrative client tries to retrieve the `ModelMBeanInfo` interface of a Version 6.1 and later MBean, the expected field names are not found in the deserialized object because of the case difference between the versions. For example, the following wsadmin function does not work when a Version 6.0.x wsadmin script connects to a Version 6.1 and later server:

```
$Help attributes MBeanObjectName
$Help operations MBeanObjectName
$Help all MBeanObjectName
```

where `MBeanObjectName` is a string representation of a Version 6.1 and later MBean Object.

To avoid this problem, set the `jmx.serial.form` Java virtual machine (JVM) custom property on the JVM custom properties page in the administrative console. Create the custom property by specifying the name value-pair on the Version 6.1 and later Application Server that you are connecting to from a Version 6.0.x client. The field names are forced to lower case to be compatible with what the Version 6.0.x client expects. The lower case field names contradict the JMX specification and compromise interoperability with future versions. Therefore, the recommendation is that you set this property only when it is absolutely needed in a mixed version environment.

Property name	<code>jmx.serial.form</code>
Data type	string
Value	1.2.0 or 1.2.1

To access the JVM custom properties page, click:

Servers > Server Types > WebSphere application servers > *server1*. Then, under Server Infrastructure, click **Java and process management > Process definition > Control > Java virtual machine > Custom properties**.

Managed object metadata

Information about a node, such as operating system platform and product features, is maintained in the configuration repository in the form of properties. As product features are installed on a node, new property settings are added.

WebSphere Application Server system management uses the managed object metadata properties as follows:

- To display the node version in the administrative console
- To ensure that new configuration types or attributes are not created or set on older release nodes
- To ensure that new resource types are not created on old release nodes
- To ensure that new applications are not installed on old release nodes because the old run time cannot support the new applications

Base properties

The following base property keys are defined for WebSphere Application Server:

com.ibm.websphere.baseProductVersion: The version of WebSphere Application Server that is installed.

com.ibm.websphere.nodeOperatingSystem: The operating system platform on which the node runs.

com.ibm.websphere.deployed.features: A list of features that extends a profile. An example of a feature is an administrative console plug-in.

com.ibm.websphere.nodeSysplexName: The sysplex name on a z/OS operating system.

Here are examples of metadata property values. The `com.ibm.websphere.deployed.features` metadata property value is split on multiple lines for printing purposes.

```
com.ibm.websphere.baseProductVersion=7.0.0.0
com.ibm.websphere.deployed.features=
com.ibm.ws.base_6.0.0.0,com.ibm.ws.j2ee_6.0.0.0,
com.ibm.ws.uddi_6.0.0.0,com.ibm.ws.wsgateway_6.0.0.0
com.ibm.websphere.nodeSysplexName=PLEX1
com.ibm.websphere.nodeOperatingSystem=os390
```

For detailed information on metadata properties, view the `ManagedObjectMetadataHelper` class in the application programming interfaces documentation.

Accessing managed object metadata properties

An administrator can query managed object metadata through the `wsadmin` tool or Application Server APIs. They can additionally be viewed on the Node Installation properties administrative console panel. This article provides details on the Application Server API method.

An accessor class is used to obtain the managed object metadata properties. An accessor instance is created through its factory. A helper class, which uses the accessor instance, makes it easy to query the base metadata properties. These classes are all part of the `com.ibm.websphere.management.metadata` package in the Application Server API documentation. The specific names of these classes are:

- `com.ibm.websphere.management.metadata.ManagedObjectMetadataHelper`
- `com.ibm.websphere.management.metadata.ManagedObjectMetadataAccessor`
- `com.ibm.websphere.management.metadata.ManagedObjectMetadataAccessorFactory`

Using command-line tools

The product provides many tools that you can call from a command line.

About this task

There are several command-line tools that you can use to start, stop, and monitor application server processes and nodes. These tools only work on local servers and nodes. They cannot operate on a remote server or node. To administer a remote server, you can use the wsadmin scripting program connected to the deployment manager for the cell in which the target server or node is configured.

The following information is common for each command:

- Unless otherwise specified, all of the application server commands are located in the *app_server_root/bin* directory.
- Parameter values that specify a server name, a node name or a cell name are case sensitive. For example, if you want to start the application server MyServer for the profile test, invoke `startServer MyServer -profileName test`. If you specify `myserver` for the server name, the `startServer` script fails.

Use the following general steps to run a command:

Procedure

1. Open a system command prompt.
2. Determine whether to run the script from the profile or application server root directory.

Most command-line tools function relative to a particular profile. To determine if a command requires the `-profileName` parameter, refer to the documentation for that specific command. If you run a command from the *app_server_root/bin* directory and do not specify the `-profileName` parameter, the default profile for the product sets profile-specific variables. To specify a different profile, use one of the following options:

- Navigate to the *app_server_root/bin* directory and run the following command, specifying the profile of interest as the value for the `profileName` parameter:

```
startServer server1 -profileName AppServerProfile
```
- When a profile is created, the application server creates a proxy script in the *profile_root/bin* directory for each script in the *app_server_root/bin* directory that is applicable to the type of profile created. When a proxy script is invoked, the profile-specific variables for the script are set based on the profile from which the script is invoked. To run the command for a specific profile, navigate to the *profile_root/bin* directory for the profile of interest to run the command.

3. Run the command of interest.

Results

The command runs the requested function and displays the results on the screen.

Refer to the command log file for additional information. When you use the `-trace` option for the command, the additional trace data is captured in the command log file. The directory location for the log files is under the default system log root directory, except for commands related to a specific server instance, in which case the log directory for that server is used. You can override the default location for the command log file using the `-logfile` option for the command.

What to do next

startServer command

The `startServer` command reads the configuration file for the specified server process and starts that server process.

Use the `startServer` command to:

- Start an application server.

- Start a DMZ Secure Proxy Server for IBM WebSphere Application Server.
- Start an administrative agent.
- Start a job manager.

Note: This topic references one or more of the application server log files. Beginning in WebSphere Application Server Version 8.0 you can configure the server to use the High Performance Extensible Logging (HPEL) log and trace infrastructure instead of using `SystemOut.log`, `SystemErr.log`, `trace.log`, and `activity.log` files or native z/OS logging facilities. If you are using HPEL, you can access all of your log and trace information using the LogViewer command-line tool from your server profile bin directory. See the information about using HPEL to troubleshoot applications for more information on using HPEL.

You do not have to use a user name and password with the `startServer` command because this command launches a server process but does not invoke an MBean method.

gotcha: You can use the administrative console to change the Java virtual machine Classpath setting or the environment entries settings for a server. However, before making these changes you should understand the following consequences of making these changes:

- If you change the value of the Java virtual machine Classpath setting, then this new value overrides the value of the Classpath parameter in the launch command that is set, by default, in the script that is generated when the you issue the `startServer -script` command.
- If you add a new environment entry on the Environment entries page or change the setting of an existing entry, then the new and changed values appear as parameters in the script that is generated when you issue the `startServer -script` command.
- If one of the environment entries you add is called PATH, then the value specified for this entry overrides the value specified for the PATH variable that, by default, is set to `WAS_PATH` in the `setUpCmdLine` file. If the value of the PATH variable is overridden, the following message is sent to the file where your error messages are logged:

```
WSVR0009E: Error occurred during startup. com.ibm.ws.exception.RuntimeError:
java.lang.NoClassDefFoundError: com/ibm/ws/process/Win32ProcessG1ue
```

For more information about where to run this command, see the *Using command line tool* topic.

Syntax

The command syntax is one of the following:

```
startServer server_name [options]
```

where *server_name* is the name of the application server or the DMZ Secure Proxy Server for IBM WebSphere Application Server that you want to start.

This argument is required.

```
startServer <adminagent_name>
```

where *adminagent_name* is the name of the administrative agent that you want to start.

```
startServer <job_manager>
```

where *job_manager* is the name of job manager that you want to start.

Parameters

The following options are available for the **startServer** command:

-? Prints a usage statement.

-help

Prints a usage statement.

-J <java_option>

Specifies options to pass through to the Java interpreter.

-logfile <fileName>

Specifies the location of the log file to which trace information is written. By default, the log file is named `startServer.log` and is created in your `logs` directory.

-profileName

Defines the profile of the server process in a multi-profile installation. The `-profileName` option is not required for running in a single profile environment. The default for this option is the default profile.

-quiet

Suppresses the progress information that the **startServer** command prints in normal mode.

-replacelog

Replaces the log file instead of appending to the current log.

-recovery

Specifies that the server will start in recovery mode, perform a transactional recovery, and shut down. The server will not accept any new transactions while it is in recovery mode. When you start the server again, resources that were unavailable due to questionable transactions will be available.

Use this option if a server fails and you do not want to accept new transactions during the recovery process.

transition: If you are migrating from a previous version of the product, make sure that the ENV parameter included on the JCL procedure statement for the controller includes either the REC=N or the REC=Y element. If the ENV parameter does not include either the REC=N or the REC=Y element, the server will not restart in recovery mode even if you specify the `-recovery` option.

If the ENV parameter includes the REC=N element, the setting is automatically changed to REC=Y if you specify `-recovery` when you restart the server. The REC=N element is automatically included on the ENV parameter if you did not migrate from a previous version of the product. Following is an example of what your updated PROC statement might look like:

```
//BB06ACR PROC ENV=,PARMS=' ',REC=N,Z=BB06ACRZ
```

-statusport <portNumber>

An optional parameter that allows an administrator to set the port number for server status callback. The tool opens this port and waits for status callback from the server indicating that the server has started. If the parameter is not set, an unused port is automatically allocated.

-script [<script fileName>] -background

Generates a launch script with the `startServer` command instead of launching the server process directly. The launch script name is an optional argument. If you do not supply the launch script name, the default script file name is `start_server` based on the server name that is passed as the first argument to the **startServer** command. The `-background` parameter is an optional parameter that specifies that the generated script will run in the background when you run it.

-trace

Generates trace information to the log file for debugging purposes.

-timeout <seconds>

Specifies the waiting time before server initialization times out and returns an error.

Usage scenario

The following examples demonstrate correct syntax. The information within the parentheses is a description of the output that is created if you issue the preceding command.

```
startServer server1
startServer server1 -script (produces the start_server1.sh file)
startServer server1 -trace (produces the startserver.log file)
startServer adminagent
startServer jobmgr
startServer proxy1 (for the secure proxy profile)
startServer.sh server1 -trace -username MyUserName -password MyUserPassword
-profileName MyProfileName (starts the server1 server using the
MyProfileName profile. The server runs under the user name MyUserName,
and produces trace files under the profile_root/logs directory)
```

stopServer command

The stopServer command reads the configuration file for the specified server process. This command sends a Java management extensions (JMX) command to the server telling it to shut down.

The server process can be an application server, a DMZ Secure Proxy Server for IBM WebSphere Application Server, an administrative agent server, or a job manager server

By default, the stopServer command does not return control to the command line until the server completes the shutdown process. There is a -nowait option to return immediately, and other options to control the behavior of the stopServer command. For more information about where to run this command, see the *Using command-line tools* topic.

gotcha: If message ADMC0074E: Connection will be closed due to unrecoverable error is displayed in the system log file, a client might have attempted to send a user name, and password to another server on which security is disabled. You can ignore this message because the client automatically changes its security setting to match the security setting for the server with which the client is trying to communicate. The connection is eventually successfully completed. However, security is now disabled for both the administrative client and the receiving server.

Syntax

The command syntax is one of the following:

```
stopServer <server_name> [options]
```

where *server_name* is the name of the configuration directory of the application server or the DMZ Secure Proxy Server for IBM WebSphere Application Server that you want to stop.

This argument is required.

```
stopServer <adminagent_name>
```

where *adminagent_name* is the name of the administrative agent that you want to stop.

```
stopServer <job_manager>
```

where *job_manager* is the name of job manager that you want to stop.

Parameters

The following options are available for the stopServer command:

- nowait**
Tells the **stopServer** command not to wait for successful shutdown of the server process.
- quiet**
Suppresses the progress information that the **stopServer** command prints in normal mode.
- logfile <fileName>**
Specifies the location of the log file to which trace information is written. By default, the log file is named `stopServer.log` and is created in the `logs` directory.
- profileName**
Defines the profile of the server process in a multi-profile installation. The `-profileName` option is not required for running in a single profile environment. The default for this option is the default profile.
- replacelog**
Replaces the log file instead of appending to the current log.
- trace**
Generates trace information into a file for debugging purposes. The trace output is written to the `stopServer.log` file which is located in the `profile_root/logs/server` directory.
- timeout <seconds>**
Specifies the time to wait for server shutdown before timing out and returning an error.
- statusport <portNumber>**
An optional parameter that allows an administrator to set the port number for server status callback. The tool opens this port and waits for status callback from the server just before the server has stopped. If the parameter is not set, an unused port is automatically allocated.
- conntype <type>**
Specifies the JMX connector type to use for connecting to the deployment manager. Valid types are SOAP or Remote Method Invocation (RMI).
- port <portNumber>**
Specifies the server JMX port to use explicitly, so that you can avoid reading the configuration files to obtain the information.
- username <name>**
Specifies the user name for authentication if security is enabled in the server. Acts the same as the `-user` option.
- user <name>**
Specifies the user name for authentication if security is enabled in the server. Acts the same as the `-username` option.
- password <password>**
Specifies the password for authentication if security is enabled in the server.

gotcha: If you are running in a secure environment but have not provided a user ID and password, you receive the following error message:

```
ADMN0022E: Access denied for the stop operation on Server MBean due to insufficient or empty credentials.
```


To solve this problem, provide the user ID and password information.
- help**
Prints a usage statement.
- ?** Prints a usage statement.

Usage scenario

The following examples demonstrate correct syntax:

```
stopServer server1
stopServer server1 -nowait
stopServer server1 -trace (produces the stopserver.log file)
stopServer adminagent
stopServer jobmgr
stopServer proxy1 (for the secure proxy profile)
```

startManager command

Use the `startManager` command to manipulate a deployment manager with scripting.

The **startManager** command reads the configuration file for the WebSphere Application Server, Network Deployment manager process and constructs a **launch** command. Depending on the options you specify, the **startManager** command launches a new Java virtual machine (JVM) API to run the manager process, or writes the **launch** command data to a file.

You must run this command from the `profile_root/bin` directory of a WebSphere Application Server, Network Deployment installation.

You do not have to use a user name and password with the **startManager** command because this command launches a server process but does not invoke an MBean method.

For more information about where to run this command, see Using command line tools.

Syntax

The command syntax is as follows:

```
startManager [options]
```

Parameters

The following options are available for the **startManager** command:

-quiet

Suppresses the progress information that the **startManager** command prints in normal mode.

-logfile <fileName>

Specifies the location of the log file to which trace information is written. By default, the log file is named `startServer.log` and is created in your `logs` directory.

-replacelog

Replaces the log file instead of appending to the current log.

-trace

Generates trace information into a file using the **startManager** command for debugging purposes.

-timeout <seconds>

Specifies the waiting time before deployment manager initialization times out and returns an error.

-statusport <portNumber>

An optional parameter that allows an administrator to set the port number for deployment manager status callback. The tool opens this port and waits for status callback from the deployment manager indicating that the deployment manager has started. If the parameter is not set, an unused port is automatically allocated.

-script [<script fileName>] -background

Generates a launch script with the **startManager** command instead of launching the deployment manager process directly. The launch script name is an optional argument. If you do not provide the launch script name, the default script file name is <start_dmgr>. The -background parameter is an optional parameter that specifies that the generated script will run in the background when you execute it.

-help

Prints a usage statement.

-? Prints a usage statement.

Usage scenario

The following examples demonstrate correct syntax:

```
startManager
```

```
startManager -script (produces the start_dmgr.sh file)
```

```
startManager -trace (produces the startmanager.log file)
```

stopManager command

The **stopManager** command reads the configuration file for the WebSphere Application Server, Network Deployment manager process.

It sends a Java Management Extensions (JMX) command to the manager telling it to shut down. By default, the **stopManager** command waits for the manager to complete the shutdown process before it returns control to the command line. There is a -nowait option to return immediately, as well as other options to control the behavior of the **stopManager** command. For more information about where to run this command, see the Using command tools article.

Syntax

The command syntax is as follows:

```
stopManager [options]
```

Parameters

The following options are available for the **stopManager** command:

-nowait

Tells the **stopManager** command not to wait for successful shutdown of the deployment manager process.

-quiet

Suppresses the progress information that the **stopManager** command prints in normal mode.

-logfile <fileName>

Specifies the location of the log file to which trace information is written. By default, the log file is named stopServer.log and is created in your logs directory.

-replacelog

Replaces the log file instead of appending to the current log.

-trace

Generates trace information to a file for debugging purposes. The trace output is written to the stopmanager.log file which is located in the *profile_root*/logs directory.

-timeout <seconds>

Specifies the waiting time for the manager to complete shutdown before timing out and returning an error.

-statusport <portNumber>

An optional parameter that allows an administrator to set the port number for deployment manager status callback. The tool opens this port and waits for status callback from the deployment manager just before the deployment manager has completely stopped. If the parameter is not set, an unused port is automatically allocated.

-conntype <type>

Specifies the Java Management Extensions (JMX) connector type to use for connecting to the deployment manager. Valid types are SOAP or Remote Method Invocation (RMI).

-port <portNumber>

Specifies the deployment manager JMX port to use explicitly, so that you can avoid reading the configuration files to obtain information.

-username <name>

Specifies the user name for authentication if security is enabled in the deployment manager. Acts the same as the `-user` option.

-user <name>

Specifies the user name for authentication if security is enabled in the deployment manager. Acts the same as the `-username` option.

-password <password>

Specifies the password for authentication if security is enabled in the deployment manager.

Note: If you are running in a secure environment but have not provided a user ID and password, you receive the following error message:

```
ADMN0022E: Access denied for the stop operation on Server MBean due
to insufficient or empty credentials.
```

To solve this problem, provide the user ID and password information.

-help

Prints a usage statement.

-? Prints a usage statement.

Usage scenario

The following examples demonstrate correct syntax:

```
stopManager
```

```
stopManager -nowait
```

```
stopManager -trace (produces the stopmanager.log file)
```

startNode command

The **startNode** command reads the configuration file for the node agent process and constructs a **launch** command.

You do not have to use a user name and password with the **startNode** command because this command launches a server process but does not invoke an MBean method.

For more information about where to run this command, see the Using command line tools article.

Syntax

The command syntax is as follows:

```
startNode [options]
```

Parameters

The following options are available for the **startNode** command:

-? Prints a usage statement.

-help

Prints a usage statement.

-logfile <fileName>

Specifies the location of the log file to which trace information is written. By default, the log file is named `startServer.log` and is created in your `profile_root/logs` directory.

-quiet

Suppresses the progress information that the **startNode** command prints in normal mode.

-replacelog

Replaces the log file instead of appending to the current log.

-statusport <portNumber>

An optional parameter that allows an administrator to set the port number for node agent status callback. The tool opens this port and waits for status callback from the node agent indicating that the node agent has started. If the parameter is not set, an unused port is automatically allocated.

-trace

Generates trace information to the log file for debugging purposes.

-timeout <seconds>

Specifies the waiting time before node initialization times out and returns an error.

Usage scenario

The following examples demonstrate correct syntax:

```
startNode
```

```
startNode -trace (produces additional trace statements in the startnode.log file)
```

stopNode command

The **stopNode** command reads the configuration file for the node agent process and sends a Java Management Extensions (JMX) command telling the node agent to shut down.

By default, the **stopNode** command waits for the node agent to complete shutdown before it returns control to the command line. There is a `-nowait` option to return immediately, as well as other options to control the behavior of the **stopNode** command. For more information about where to run this command, see the Using command tools topic.

If you stop the server before stopping the node agent using the **stopserver -servers** command, the server will not restart when you issue a **startserver** command.

Syntax

The command syntax is as follows:

```
stopNode [options]
```

Parameters

The following options are available for the **stopNode** command:

-nowait

Tells the **stopNode** command not to wait for successful shutdown of the node agent process.

-quiet

Suppresses the progress information that the **stopNode** command prints in normal mode.

-logfile <fileName>

Specifies the location of the log file to which trace information is written. By default, the log file is named `stopServer.log` and is created in the `logs` directory of the profile for the node being stopped.

-replacelog

Replaces the log file instead of appending to the current log.

-trace

Generates trace information into a file for debugging purposes. The trace output is written to the `was_jobname-jobuser-jobnum.log` file which is located in the `profile_root/logs` directory.

-timeout <seconds>

Specifies the waiting time for the agent to shut down before timing out and returning an error.

Note: The **timeout** parameter is used to stop one server at a time. If you use the **timeout** parameter with the **stopservers** parameter, the **timeout** parameter value is ignored.

-statusport <portNumber>

An optional parameter that allows an administrator to set the port number for node agent status callback. The tool opens this port and waits for status callback from the node agent just before the node agent has completely stopped. If the parameter is not set, an unused port is automatically allocated.

-stopservers [-saveNodeState]

Stops all application servers on the node before stopping the node agent.

Use the `-saveNodeState` option to save the server restart state so that you can start the application servers based on the previous state using the **startNode** command.

-conntype <type>

Specifies the Java Management Extensions (JMX) connector type to use for connecting to the deployment manager. Valid types are SOAP or Remote Method Invocation (RMI).

-port <portNumber>

Specifies the node agent JMX port to use explicitly, so that you can avoid reading configuration files to obtain the information.

-username <name>

Specifies the user name for authentication if security is enabled in the node agent. Acts the same as the `-user` option.

-user <name>

Specifies the user name for authentication if security is enabled in the node agent. Acts the same as the `-username` option.

-password <password>

Specifies the password for authentication if security is enabled in the node agent.

Note: If you are running in a secure environment but have not provided a user ID and password, you receive the following error message:

ADMN0022E: Access denied for the stop operation on Server MBean due to insufficient or empty credentials.

To solve this problem, provide the user ID and password information.

-help

Prints a usage statement.

Note: When requesting help for the usage statement for the **stopNode** command, a reference to the **stopServer** command displays. All of the options displayed for this usage statement apply to the **stopNode** command.

-? Prints a usage statement.

Note: When requesting help for the usage statement for the **stopNode** command, a reference to the **stopServer** command displays. All of the options displayed for this usage statement apply to the **stopNode** command.

Usage scenario

The following examples demonstrate correct syntax:

```
stopNode
```

```
stopNode -nowait
```

```
stopNode -trace (produces the stopnode.log file)
```

serverStatus command

Use the **serverStatus** command to obtain the status of one or all of the servers configured on a node.

For more information about where to run this command, see the *Using command line tools* topic.

Syntax

The command syntax is as follows:

```
serverStatus <server>|-all [options]
```

The first argument is required. The argument is either the name of the server for which status is desired, or the **-all** keyword which requests status for all servers defined on the node.

Parameters

The following options are available for the **serverStatus** command:

-quiet

Suppresses the progress information that the **serverStatus** command prints in normal mode.

-logfile <fileName>

Specifies the location of the log file to which trace information is written. By default, the log file is named `serverStatus.log` and is created in your `logs` directory.

-profileName

Defines the profile of the Application Server process in a multi-profile installation. The **-profileName** option is not required for running in a single profile environment. The default for this option is the default profile.

-replaceLog

Replaces the log file instead of appending to the current log.

-trace

Generates trace information into a file for debugging purposes.

- username <name>**
Specifies the user name for authentication if security is enabled. Acts the same as the -user option.
- user <name>**
Specifies the user name for authentication if security is enabled. Acts the same as the -username option.
- password <password>**
Specifies the password for authentication if security is enabled.
- help**
Prints a usage statement.
- ?** Prints a usage statement.

Usage scenario

The following examples demonstrate correct syntax:

```
serverStatus server1
```

```
serverStatus -all (returns status for all defined servers)
```

```
serverStatus -trace (produces the serverStatus.log file)
```

addNode command

The addNode command incorporates an application server installation into a cell.

Depending on the size and location of the new node you incorporate into the cell, this command can take a few minutes to complete.

You must have Administrator privileges to use the addNode function.

The node agent server is automatically started as part of the addNode command unless you specify the -noagent option. If you recycle the system that hosts an application server node, and did not set up the node agent to act as an operating system daemon, you must issue a startNode command to start the node agent before starting any application servers.

When you run the addNode command, the command stops the running application server that it is incorporating into a cell. You can optionally stop the application server before running the addNode command.

When you add a node, the product might generate ports. The following items apply to port generation:

- Ports generated for the node agent are unique for all the profiles in the installation. For development purposes, you can create multiple profiles on the same installation and add them to one or more cells without concern for port conflicts.
- If you want to specify the ports that the node agent uses, specify the ports in a file with the file name passed with the -portprops option. The format of the file is key=value pairs, one on each line, with the key being the same as the port name in the serverindex.xml file.
- If you want to use a number of sequential ports, consider the -startingport option. This means that port conflicts with other profiles are not detected.

Note: Use the -includeapps option for the addNode command to ensure that the environment starts with the same version of the application. If any custom policy set is created on the server before you run the addNode command, then the custom policy set is not copied to the new cell after you run the addNode command. Therefore, when using the -installApps option, an application on the server that uses the custom policy set fails to load the policy set after you run the addNode command. You can

export the custom policy set from the stand-alone server before you run the addNode command, and import the custom policy set to the new cell after you run the addNode command.

Read the topic on using command-line tools to determine whether to run the command from the profile or application server root directory.

Syntax

See the command syntax:

```
addNode dmgr_host [dmgr_port] [-conntype type] [-includeapps] [-includebuses]
[-startingport portnumber] [-portprops qualified_filename]
[-nodeagentshortname name] [-nodegroupname name] [-registerservice]
[-serviceusername name] [-servicepassword password] [-coregroupname name]
[-noagent] [-statusport 1231] [-quiet] [-nowait] [-logfile filename]
[-replacelog] [-trace] [-username uid] [-password pwd]
[-localusername localuid] [-localpassword localpwd] [-profileName profilename]
[-excludesecuritydomains true | false] [-asExistingNode] [-help]
```

The dmgr_host argument is required. All the other arguments are optional. The default port number is 8879 for the default SOAP port of the deployment manager. SOAP is the default Java Management Extensions (JMX) connector type for the command. If you have multiple product installations or multiple profiles, the SOAP port might be different from 8879. Examine the deployment manager SystemOut.log file to see the current ports in use.

Note: This topic references one or more of the application server log files. Beginning in WebSphere Application Server Version 8.0 you can configure the server to use the High Performance Extensible Logging (HPEL) log and trace infrastructure instead of using SystemOut.log, SystemErr.log, trace.log, and activity.log files or native z/OS logging facilities. If you are using HPEL, you can access all of your log and trace information using the LogViewer command-line tool from your server profile bin directory. See the information about using HPEL to troubleshoot applications for more information on using HPEL.

Parameters

The following options are available for the addNode command:

-conntype <type>

Specifies the JMX connector type to use for connecting to the deployment manager. SOAP is the default Java Management Extensions (JMX) connector type for the command. Other valid types are JSR160RMI or Remote Method Invocation (RMI).

-includeapps

By default the addNode command does not carry over applications from the stand-alone servers on the new node to the cell. In general, install applications using the deployment manager. The -includeapps option tells the addNode command to carry over the applications from a node. If the application already exists in the cell, then a warning is printed and the application does not install in the cell.

The applications are mapped to the server that you federated using the addNode command. When the addNode command operation completes, the applications run on that server when the server is started. Since these applications are part of the network deployment cell, you can map them to other servers and clusters in the cell using the administrative console. Read about mapping modules to servers in the *Developing and deploying applications* PDF for more information.

Do not use the -includeapps option if the node that you want to federate includes the product-supplied applications, such as the Samples. If you do, the second node to be federated that includes these applications is rejected because the applications exist in the cell and application merge is not supported.

If you use the -includeapps option on a node that includes a large number of applications, then the timeout for the administrative connector must be extended to account for the additional time required

to transfer all the applications to the deployment manager during the `addNode` operation and to remotely install them into the cell. The `-includeapps` option is not a recommended approach unless only a few unique applications exist on the node.

By default, during application installation, application binaries are extracted in the `app_server_root/installedApps/cellName` directory. After the `addNode` command, the cell name of the configuration on the node that you added changes from the base cell name to the deployment manager cell name. The application binary files are located where they were before you ran the `addNode` command, for example, `app_server_root/installedApps/old_cellName`.

In the following example the application was installed by explicitly specifying the location for binary files:

```
${APP_INSTALL_ROOT}/${CELL}
```

The variable `${CELL}`, specifies the current cell name. Then, when the `addNode` command runs, the binary files are moved to the following directory:

```
profile_root/installedApps/currentCellName
```

Federating the node to a cell using the `addNode` command does not merge any cell-level configuration, including virtual host information. If the virtual host and aliases for the new cell do not match the product, you cannot access the applications running on the servers. You have to manually add all the virtual host and host aliases to the new cell, using the administrative console running on the deployment manager.

Note: When the `-includeapps` parameter is specified, an `OutOfMemoryError` might occur if the Java virtual machine (JVM) heap size is too small. When this error occurs, the following error message is issued:

```
ADMU0111E: Program exiting with error: java.lang.OutOfMemoryError
```

This error can occur when large applications are processed, or when there is a large number of applications in the Base Application Server.

To recover from this error and successfully federate the application server, complete the following actions:

1. Issue the `cleanupNode` command on your deployment manager server. Read about the `cleanupNode` command for more information about this command.
2. Increase the JVM heap size for the `addNode` script. When you issue the `addNode` command, the JVM heap size is set to `-Xms128m -Xmx512m`. To increase these values, use the `-javaoption` parameter. For example, you might specify the following (all on one line):
3. Reissue the `addNode` command.

-includebuses

Copies the buses from the node to be federated to the cell. This parameter also attempts to copy the service integration bus configuration of the remote node into the cell. If the destination cell already contains a bus with the same name as any bus at the remote node, the `addNode` command fails. To prevent this failure, you can act before using the `addNode` command. You can delete the bus with that name in the destination cell, rename the bus to be added to the cell, or manually configure the bus already located in the cell.

-startingport <portNumber>

Supports the specification of a port number to use as the base port number for all node agent and Java Messaging Service (JMS) server ports created when the `addNode` command runs. With this support you can control which ports are defined for these servers, rather than using the default port values. The starting port number is incremented one unit to calculate the port number for every node agent port and JMS server port configured during the `addNode` command.

If multiple node agents exist on the same physical server, then you can define the base port number for each by using the `-startingport` parameter before federation, or by modifying the ports in the node agent section of the `serverindex.xml` file.

-portprops <filename>

Passes the name of the file that contains key-value pairs of explicit ports that you want the new node agent to use. For example, to set your SOAP and RMI ports to 3000 and 3001, create a file with the following two lines and pass it as the parameter:

```
SOAP_CONNECTOR_ADDRESS=3000
BOOTSTRAP_ADDRESS=3001
```

-nodeagentshortname <name>

The shortname to use for the new node agent.

-nodegroupname <name>

The name of the node group in which to add this node. If you do not specify, the node is added to the `DefaultNodeGroup`.

Registers the node agent as a Windows service.

You can optionally define a user name and password for the windows service using the `-serviceusername` parameter and the `-servicepassword` parameter. If you define a user name, you must give the user name `Logon` as a service authority for the service to run properly.

If you do not specify a user name and password, then the service runs under the local system account.

-coregroupname <name>

The name of the core group in which to add this node. If you do not specify this option, the node is added to the `DefaultCoreGroup`.

-noagent

Tells the `addNode` command not to launch the node agent process for the new node.

-statusport

An optional parameter that allows an administrator to set the port number for node agent status callback. The tool opens this port and waits for status callback from the node agent indicating that the node agent has started. If the parameter is not set, an unused port is automatically allocated.

-quiet

Suppresses the progress information that the `addNode` command prints in normal mode.

-nowait

Tells the `addNode` command not to wait for successful initialization of the launched node agent process.

-logfile <filename>

Specifies the location of the log file to which trace information is written. By default, the log file is `addNode.log` and is created in the logs directory of the profile for the node being added.

-replacelog

Replaces the log file instead of appending to the current log file. By default, the `addNode` command appends to the existing trace file. This option causes the `addNode` command to overwrite the trace file.

-trace

Generates additional trace information in the log file for debugging purposes.

-user <name> or -username <name>

Specifies the user name for authentication if security is enabled. Acts the same as the `-user` option. The user name that you choose must be a pre-existing user name.

-password <password>

Specifies the password for authentication if security is enabled. The password that you choose must be one that is associated with a pre-existing user name.

-localusername <name>

Specifies the user name for authentication for existing application servers on the node that you want to federate. This parameter is only applicable if security is enabled for the application server.

-localpassword <password>

Specifies the password for authentication for existing application servers on the node that you want to federate. The password that you choose must be one that is associated with a pre-existing user name. This parameter is only applicable if security is enabled for the application server.

-profileName

Defines the profile of the Application Server process in a multi-profile installation. The `-profileName` option is not required for running in a single profile environment. The default for this option is the default profile. If you are adding a non-default profile to the deployment manager cell, then this parameter is required.

-excludesecuritydomains true | false

Set the `-excludesecuritydomains` parameter to true if you do not want the security domains configured at the application server node federated into the cell. When the parameter is set to true, the security configuration of the cell is used. This parameter applies only when you have security domains configured at the unfederated application server. By default, if there is a security domain associated with an application server, the security domain is federated to the cell so that the server uses the same security domain information after it is federated.

-asExistingNode

Specifies to recover an existing managed node of a deployment manager cell.

Use the `-asExistingNode` parameter of the `addNode` command to quickly recover a damaged node. For example, if a machine failure results in an unavailable node but node information remains on the deployment manager, you can use the `addNode -asExistingNode` option to recreate the unavailable node by completing the following steps:

1. Create a new profile with the same node and profile name as the unavailable node. You can create the profile on a different machine from the original node.
2. For the new profile, run the `addNode` command with the `-asExistingNode` option.

You can also use the `-asExistingNode` option of the `addNode` command to move a node to a product installation on a different computer but at the same path, to move a node to a product installation on a different operating system or with a different path, or to create cells from a template cell. See the topic on recovering or moving nodes with the `addNode -asExistingNode` command.

Note: Other `addNode` options for node configuration are incompatible with this `-asExistingNode` option. Do not use `-asExistingNode` with the following incompatible options: `-includeapps`, `-includebuses`, `-startingport`, `-portprops`, `-nodeagentshortname`, `-nodegroupname`, `-registerservice`, `-serviceusername`, `-servicepassword`, `-coregroupname`, or `-excludesecuritydomains`.

-help

Prints a usage statement.

`-?` Prints a usage statement.

Usage scenario

The following examples demonstrate correct syntax:

```
addNode testhost 8879 (adds an application server to the deployment manager)
```

```
addNode deploymgr 8879 -trace (produces the addNode.log file)
```

```
addNode host25 8879 -nowait (does not wait for a node agent process)
```

The value 8879 is the default port.

Security considerations when adding an application server node to WebSphere Application Server, Network Deployment cell

When adding a node to a cell, the newly federated node automatically inherits the user registry (Local OS, Lightweight Directory Access Protocol (LDAP), or Custom), authentication mechanism (LTPA), and authorization setting (WebSphere bindings or System Authorization Facility (SAF) EJBROLE profiles) of the existing WebSphere Application Server, Network Deployment cell.

For distributed security, all servers in the cell must use the same user registry and authentication mechanism. To recover from a user registry change, you must modify your applications so that the user and group-to-role mappings are correct for the new user registry. See the article on assigning users and groups to roles.

Another important consideration is the Secure Sockets Layer (SSL) public-key infrastructure. Before running the `addNode` command with the deployment manager, verify that the `addNode` command can communicate as an SSL client with the deployment manager. This communication requires that the `addNode` truststore that is configured in the `sas.client.props` file contains the signer certificate of the deployment manager personal certificate, as found in the keystore and specified in the administrative console.

The following issues require consideration when running the `addNode` command with security:

- When attempting to run system management commands such as the `addNode` command, specify administrative credentials to perform the operation. The `addNode` command accepts `-username` and `-password` parameters to specify the user ID and password. The user ID and password that are specified must be for an administrative user. For example, specify a user that is a member of the console users with Administrator privileges or the administrative user ID configured in the user registry. See the following example of the `addNode` command:

```
addNode CELL_HOST 8879 -includeapps -username user -password pass
```

The `-includeapps` parameter is optional. This option attempts to include the server applications into the Deployment Manager. The `addNode` command might fail if the user registries used by the application server and the deployment manager are not the same. To correct this failure, either make the user registries the same or turn off security. If you change the user registries, remember to verify that the users-to-roles and groups-to-roles mappings are correct. Read about the `addNode` command for more information about the `addNode` syntax.

If you issue the `addNode` command with security enabled, you must use a user ID with authority and specify the `-user` and `-password` options.

- Adding a secured remote node through the administrative console is not supported. You can either disable security on the remote node before performing the operation or perform the operation from the command prompt using the `addNode` script.
- Before running the `addNode` command, you must verify that the truststore files on the nodes communicate with the keystore files and System Authorization Facility (SAF) keyring that is owned by the deployment manager and vice versa. If you generate certificates for the deployment manager using the same certificate authority as you used for the node agent process, then you are successful. The following SSL configurations must contain keystores and truststores that can interoperate:
 - System SSL repertoire that is specified in the administrative console. Click **System administration > Deployment manager > HTTP transports > sslportno > SSL**.
 - SSL repertoire for appropriate JMX connector if SOAP is specified. Click **System administration > Deployment manager > Administration Services > JMX Connectors > SOAPConnector > Custom properties > sslConfig**.
 - SSL repertoire that is specified in NodeAgent. Click **System Administration > Node agents > NodeAgent Server > Administration Services > JMX Connectors > SOAPConnector > Custom properties > sslConfig**.

Use caution when adding a node to a deployment manager configuration that defines a different security domain.

- If security is enabled for the Version 7 or 8 deployment manager, then in order to federate a Version 6.x node the deployment manager cannot use the auto-generated internal server ID. The auto-generated internal server ID is used by default when enabling security.
- When a client from a previous release tries to use the `addNode` command to federate to a Version 7 or 8 deployment manager, the client must first obtain signers for a successful handshake. For more information about required changes before running the `addNode` command in this scenario, read about obtaining signers from a previous release in the topic on Secure installation for client retrieval, specifically the Obtaining signers for clients and servers from a previous release section. The user registry can be changed by performing one of the following actions:
 - On the administrative console, click **Global Security**. Under **Available realm definitions**, click **Configure > Server identity that is stored in repository**. Enter the user name and password and then click **Apply**.
 - (The following command is split on multiple lines for printing purposes.)

```
AdminTask.configureAdminWIMUserRegistry(['-autoGenerateServerId false -serverId testuser
-serverIdPassword testuserpwd -primaryAdminId testuser -ignoreCase true ]')
```

The server must be restarted for these changes to take effect.

- After running the `addNode` command, the application server is in a new SSL domain. It might contain SSL configurations that point to keystore and truststore files that are not prepared to interoperate with other servers in the same domain. Consider which servers are intercommunicating, and ensure that the servers are trusted within your truststore files.

addNode command best practices

Use the **addNode** command to add a stand-alone node into a cell.

The **addNode** command does the following:

- Copies the base WebSphere Application Server cell configuration to a new cell structure. This new cell structure matches the structure of deployment manager.
- Creates a new node agent definition for the node that the cell incorporates.
- Sends commands to the deployment manager to add the documents from the new node to the cell repository.
- Performs the first configuration synchronization for the new node, and verifies that this node is synchronized with the cell.
- Launches the node agent process for the new node.
- Updates the `setupCmdLine.bat` or `setupCmdline.sh` files and the `wsadmin.properties` file to point to the new cell.
- After federating the node, the **addNode** command backs up the `plugin-cfg.xml` file from the `app_server_root/config/cells` directory to the `config/backup/base/cells` directory. The **addNode** command regenerates a new `plugin-cfg.xml` file at the Deployment Manager and the `nodeSync` operation copies the files to the node level.

Tips for using the **addNode** command:

- If you add a node to a cell, the cell name of the node you are federating must be different from the name of the cell to which the node is federated. Otherwise, you receive the `ADMU0027E` message, and the `addNode` command does not add the node to the cell.
- Verify that the deployment manager and nodes are updated to the same revision level within the WebSphere Application Server. For example, a deployment manager at level 6.0.1 will be unable to federate with nodes at 6.0.2.
- Do not put WebSphere Application Server `.jar` files on the generic `CLASSPATH` variable (default class path) for the overall system.
- By default, applications that are installed on the node will not copy to the cell. If you install an application after using the **addNode** command, the application will install on the cell. By specifying the

-includeapps option, you force the **addNode** command to copy applications from the node to the cell. Applications with duplicate names will not copy to the cell.

- Cell-level documents are not merged. Any changes that you make to the stand-alone cell-level documents before using the **addNode** command must be repeated on the new cell. For example, virtual hosts.
- If you receive an OutOfMemory exception while using the **addNode** command, you may need to increase the heap size of the deployment manager. To increase the heap size of the deployment manager, adjust the Maximum heap size parameter. For example, in the administrative console, go to **System administration > Deployment manager > Java and Process Management > Process definition > Java Virtual Machine** and increase the **Maximum heap size** value.
- In some instances it may take longer than anticipated for the deployment manager to respond to the **addNode** command. The default timeout value, which determines how long the client will wait for a server response, is appropriate in the majority of cases. However, you may require more time for the server to respond under heavier processing conditions. For example, if you include the -includeapps option and have a large number of applications, or the applications are very large, the default value of 180 seconds may be insufficient. To change the default timeout value, open the file `$WAS_HOME/profiles/<profile name>/properties/soap.client.props` in any ASCII text editor and find the following line (shown here with default value of 180 seconds):

```
com.ibm.SOAP.requestTimeout=180
```

If you need to change the default you can edit this line to set the timeout to a value more appropriate for your situation (**Note:** setting the above value to 0 will disable the timeout check altogether). If the timeout value is set too high you will have to wait a long time to determine if the **addNode** command will successfully complete its request to the deployment manager. If the value is set too short the deployment manager will not have sufficient time to complete the request before the **addNode** command concludes that the deployment manager is not responding and will respond with an error. Other factors that may affect server timeouts include the processing load or excessive paging on the deployment manager and network latency. Some of these conditions may be transient.

- If you receive an **addNode** error message regarding bad clock syncs, make sure that the machine with the node to be federated is in time sync with the Deployment Manager machine to which the node is to be federated.
- If you use the **addNode** command from a node that was federated to an existing deployment manager, the deployment manager will be corrupt. You will not be able to start the second deployment manager after you stop it. This happens because the **addNode** command creates a `dmgrProfile]/config/cells/dmgrCell/nodeName` directory in the master configuration. This is an incomplete node configuration directory.

You will come into contact with the issue if you have a federated node and run the **addNode** command again for a different deployment manager. This causes the deployment manager to be corrupted and you will not be able to start the deployment manager afterwards because of the incomplete node directory.

Perform one of the following solutions to resolve this issue:

- If the deployment manager is running, you can use the **cleanupNode** command on deployment manager where the incomplete node resides.
- Manually delete the directory that was created on the deployment manager configuration during an **addNode** command operation that was incomplete. For example:

```
app_server_root/profiles/dmgrProfile/config/cells/dmgrCell/]nodeName
```

removeNode command

The removeNode command returns a node from a WebSphere Application Server, Network Deployment distributed administration cell to a stand-alone application server installation.

The `removeNode` command only removes the node-specific configuration from the cell. This command does not uninstall any applications that were installed as the result of running an `addNode` command. Such applications can subsequently deploy on additional servers in the WebSphere Application Server, Network Deployment cell. As a consequence, an `addNode` command with the `-includeapps` option ran after a `removeNode` command does not move the applications into the cell because they exist from the first `addNode` command. The resulting application servers added on the node do not contain any applications. To deal with this situation, add the node and use the deployment manager to manage the applications. Add the applications to the servers on the node after the node is incorporated into the cell.

Note: If your applications were modified in a feature pack environment, uninstall and reinstall the application into the base configuration after removing the node. This process ensures that your changes are not lost when removing a node configured for the feature pack. No action is necessary if your applications were not modified.

Running the `removeNode` command completes the following actions and system conditions:

- **Actions**
 - Stops all of the running server processes in the node, including the node agent process
 - Removes the node configuration documents from the cell repository by sending commands to the deployment manager
 - Copies the original application server cell configuration that existed when the node was added and the application originally install into the active configuration
- **System Conditions**
 - The EAR files remain on the deployment manager.
 - The installed application remains on the deployment manager.
 - No changes you made after federating the node will be restored to base node.
 - You have the old base node that existed before federation. After federating, you still have the original configuration for the node, the servers, and the original applications in their original states.
 - If you federate again, you overwrite any updates on the deployment manager.

Remember:

- After the node is removed from the cell, it will have the applications as they were before the node was added to the cell. Any changes you made after federating will not be included.
- The `removeNode` command returns a node from a WebSphere Application Server, Network Deployment distributed administration cell to a stand-alone application server installation.

You must have the following privileges or authorities to use the `removeNode` command:

- You must have Administrator privileges to use the `removeNode` function.

Depending on the size and location of the new node you remove from the cell, this command can take a few minutes to complete.

Removing a node that was federated as part of cell profile creation

An application server node that is built as part of a cell profile creation does not have an original configuration. Therefore, the `removeNode` command does not restore the node to a usable base configuration. If you use the `removeNode` command on a node that was created during cell profile creation, then the command indicates that the node removal utility cannot remove the node and restore the node to a base configuration.

To successfully remove a node that was federated as part of a cell profile creation, use the Profile Management Tool to delete the profile for the node. After the profile for the node is deleted, use the `cleanupNode` command on the deployment manager to remove the node configuration from the cell repository. You can create a new profile using the Profile Management Tool.

Syntax

See the following removeNode command syntax:

```
removeNode [options]
```

All the following parameters are optional.

Parameters

The following options are available for the removeNode command:

-force

Cleans up the local node configuration regardless of whether you can reach the deployment manager for cell repository cleanup. After using the -force parameter, you might need to use the cleanupNode command on the deployment manager.

-logfile <fileName>

Specifies the location of the log file to which trace information is written. By default, the log file is named removeNode.log and is created in the logs directory of the profile for the node being removed.

-password <password>

Specifies the password for authentication if security is enabled.

-quiet

Suppresses the progress information that the removeNode command prints in normal mode.

-replacelog

Replaces the log file instead of appending to the current log file.

-statusport <portNumber>

An optional parameter that allows an administrator to set the port number for server status callback. The tool opens this port and waits for status callback from the server just before the server has completely stopped. If the parameter is not set, an unused port is automatically allocated.

-trace

Generates trace information into a file for debugging purposes.

-user <name>

Specifies the user name for authentication if security is enabled. Acts the same as the -username option.

-username <name>

Specifies the user name for authentication if security is enabled. Acts the same as the -user option.

-help

Prints a usage statement.

-? Prints a usage statement.

Usage scenario

The following examples demonstrate correct syntax:

```
removeNode -quiet
```

```
removeNode -trace (produces the removeNode.log file)
```

cleanupNode command

The **cleanupNode** command cleans up a node configuration from the cell repository.

Only use this command to clean up a node if you have a node defined in the cell configuration, but the node no longer exists. For more information about where to run this command, see the Using command tools article.

Syntax

The command syntax is as follows:

```
cleanupNode <node name> [deploymgr host] [deploymgr port] [options]
```

where the first argument is required.

Parameters

The following options are available for the **cleanupNode** command:

-quiet

Suppresses the progress information that the **cleanupNode** command prints in normal mode.

-trace

Generates trace information into a log file for debugging purposes.

syncNode command

The **syncNode** command forces a configuration synchronization to occur between the node and the deployment manager for the cell in which the node is configured.

The node agent server runs a configuration synchronization service that keeps the node configuration synchronized with the master cell configuration. If the node agent is unable to run because of a problem in the node configuration, you can use the syncNode command to perform a synchronization when the node agent is not running in order to force the node configuration back in sync with the cell configuration. If the node agent is running and you want to run the syncNode command, you must first stop the node agent.

The syncNode.log file is located in the *profile_root/logs* directory.

For more information about where to run this command, see the Using command tools topic.

Syntax

The command syntax is as follows:

```
syncNode <deploymgr host> <deploymgr port> [options]
```

- The <deploymgr host> argument is required.
- The <deploymgr port> is the Java Management Extensions (JMX) port number that corresponds to the connector type on the -conntype parameter.

Parameters

The following options are available for the **syncNode** command:

-stopservers

Tells the **syncNode** command to stop all servers on the node, including the node agent, before performing configuration synchronization with the cell.

-restart

Tells the **syncNode** command to launch the node agent process after configuration synchronization completes. If the node agent is running and you want to run the syncNode command, you must first stop the node agent.

-nowait

Tells the **syncNode** command not to wait for successful initialization of the launched node agent process.

-quiet

Suppresses the progress information that the **syncNode** command prints in normal mode.

-logfile <fileName>

Specifies the location of the log file to which trace information is written. By default, the log file is named `syncNode.log` and is created in the `logs` directory of the profile for the node that you are synchronizing.

-replacelog

Replaces the log file instead of appending to the current log.

-trace

Generates trace information into a file for debugging purposes.

-timeout <seconds>

Specifies the waiting time before node agent initialization times out and returns an error.

-username <name>

Specifies the user name for authentication if security is enabled. Acts the same as the `-user` option.

-user <name>

Specifies the user name for authentication if security is enabled. Acts the same as the `-username` option.

-password <password>

Specifies the password for authentication if security is enabled.

-conntype <type>

Specifies the Java Management Extensions (JMX) connector type to use for connecting to the deployment manager. Valid types are SOAP or Remote Method Invocation (RMI). The default type is SOAP.

-help

Prints a usage statement.

-? Prints a usage statement.

Usage scenario

The following examples demonstrate correct syntax:

```
syncNode.sh testhost 8879
```

```
syncNode.sh deploymgr 8879 -trace (produces the syncNode.log file)
```

```
syncNode.sh host25 4444 -stopservers -restart  
(assumes that the deployment manager JMX port is 4444)
```

renameNode command

A short name is assigned to a node during installation and customization. Use the **renameNode** command to modify the node name of a federated server.

Launch this command from the node whose name you want to modify. When you issue this command the following occurs:

1. Connects to the deployment manager.
2. Stops all servers.
3. Changes the node configuration on the deployment manager.

4. Synchronizes the node.

For more information about where to run this command, see the Using command line tools article.

Note: You can use the `renameNode` command to modify managed nodes. However, you cannot use it to modify the node name for a deployment manager node. If you need to modify the node name for a deployment manager node, you must modify the `WAS_NODE` variable within the `profile_root/deployment_manager_name/bin/setupCmdLine` file. For more information, see the documentation about renaming deployment manager nodes.

Syntax

The command syntax is as follows:

```
renameNode.sh dmgr_host dmgr_port node_name  
[-nodeshortname name] [-trace] [-conntype type] [-username uid]  
[-password pwd] [-logfile filename] [-help]
```

Parameters

The following options are available for the **renameNode** command:

-nodeshortname <name>

The short name of the node.

-trace

Generates additional trace information in the log file for debugging purposes.

-conntype <type>

Specifies the JMX connector type to use for connecting to the deployment manager. Valid types are SOAP or RMI, which stands for Remote Method Invocation.

-username <uid>

Specifies the user name for authentication if security is enabled. Acts the same as the `-user` option. The user name that you choose must be a pre-existing user name.

-password <pwd>

Specifies the password for authentication if security is enabled. The password that you choose must be one that is associated with a pre-existing user name.

-logfile <filename>

Specifies the location of the log file to which trace information is written. By default, the log file is called `renameNode.log` and is created in the logs directory of the profile for the node being renamed.

-help

Prints a usage statement.

-? Prints a usage statement.

Usage scenario

The following examples demonstrate correct syntax:

```
renameNode.sh localhost 8879 newnode
```

registerNode command

Use the `registerNode` command to register a stand-alone node with an administrative agent so that the administrative agent can manage the node.

Run the `registerNode` command from the `bin` directory of the administrative agent server to register a node with the administrative agent. When you run the command, the stand-alone node is converted into a node that the administrative agent manages.

The administrative agent and the node being registered must be on the same computer.

gotcha: Registered nodes must have the same products as the administrative agent, and the products must be at the same version levels on the registered node and the administrative agent. This requirement is enforced because the administrative agent must have a matching environment in order to handle all of the administrative capabilities of the registered node. A node is not allowed to register with an administrative agent unless that node has an identical set of products and versions.

You can only run the command on an unfederated node. If the command is run on a federated node, the command exits with an error.

When you run the `registerNode` command, the command stops all running application servers on the node. You can optionally stop application servers on the node that you are registering before running the `registerNode` command.

transition: If you were previously running on Version 7.0.0.11 or earlier, and have an administrative agent with a managed node that has mismatched products or versions, when you when you migrate to Version 8.0, that administrative agent will not be able to start the subsystem for any mismatched nodes. You must update these nodes to have the same products and versions as the administrative agents, restart the servers on the node and then restart the administrative agent, before the administrative agent can resume managing these registered nodes

If the administrative console or the management Enterprise JavaBeans (EJB) applications of the application server being registered are enabled, the node registration process disables them.

Syntax

The `registerNode` command syntax is as follows:

```
registerNode [options]
```

Parameters

The following options are available for the `registerNode` command:

-conntype<JSR160RMI | IPC | RMI | SOAP>

The optional connector type used to connect to the administrative agent to initiate node registration. The default is SOAP.

Note: You should eventually switch from the RMI connector to the JSR160RMI connector because support for the RMI connector is deprecated.

-host *host_name*

An optional parameter that specifies the host name of the administrative agent.

-name *managed_node_name*

An optional parameter that specifies the name of the managed node after the node is registered with the administrative agent.

-nodepassword *node_password*

An optional parameter that specifies the password of the node that you are registering. Specify this parameter if security is on at the node and the password is different from the administrative agent password. Use this parameter with the `-nodeusername` parameter. The `-nodeusername` and `-nodepassword` parameters are used to stop all servers on the node.

-nodeusername *node_user_name*

An optional parameter that specifies the user name of the node that you are registering. Specify this parameter if security is on at the node and the user name is different from the administrative agent user name. Use this parameter with the `-nodepassword` parameter. The `-nodeusername` and `-nodepassword` parameters are used to stop all servers on the node.

-openConnectors *connectors*

An optional parameter that specifies a list of connectors separated by commas of connectors that the administrative agent will open. By default, all connectors are opened.

-port *port_number*

An optional parameter that specifies the port number of the administrative agent connector port.

The default port number is 8878 for the default SOAP port of the administrative agent. SOAP is the default Java Management Extensions (JMX) connector type for the command. If you have multiple product installations or multiple profiles, the SOAP port might be different from 8878. Examine the administrative agent `SystemOut.log` file to see the current ports in use.

Note: This topic references one or more of the application server log files. Beginning in WebSphere Application Server Version 8.0 you can configure the server to use the High Performance Extensible Logging (HPEL) log and trace infrastructure instead of using `SystemOut.log`, `SystemErr.log`, `trace.log`, and `activity.log` files or native z/OS logging facilities. If you are using HPEL, you can access all of your log and trace information using the LogViewer command-line tool from your server profile bin directory. See the information about using HPEL to troubleshoot applications for more information on using HPEL.

-profilePath *profile_path*

A required parameter that specifies the path to the profile of the application server node to be registered.

-trace

An optional parameter that provides tracing output for the `registerNode` command.

-portsFile *ports_file*

An optional parameter that specifies the path to a file that defines port settings for the newly registered node.

During node registration, the `registerNode` command uses an automatically generated set of recommended ports if you do not specify the `-portsFile` parameter. The recommended port values can be different than the default port values based on the availability of the default ports.

The format of the *ports_file* file that you specify is the same as the `portdef.props` file, except that only the following ports are used: `SOAP_CONNECTOR_ADDRESS`, `RMI_CONNECTOR_ADDRESS`, `JSR160RMI_CONNECTOR_ADDRESS`, and `IPC_CONNECTOR_ADDRESS`.

-profileName *profile_name*

An optional parameter that specifies the profile name of the administrative agent.

-username *username*

An optional parameter that specifies the user ID to log on to the administrative agent.

-password *password*

An optional parameter that specifies the password to log on to the administrative agent.

-help

An optional parameter that prints a usage statement.

-? An optional parameter that prints a usage statement.

Usage scenarios

The following examples demonstrate correct syntax. Commands are split on multiple lines for printing purposes.

```
registerNode.sh -conntype SOAP -port 8878  
-profilePath app_server_root/profiles/default
```

deregisterNode command

Use the `deregisterNode` command to unregister a node from an administrative agent so that you can use the node stand-alone, register the node with another administrative agent, or federate the node with the deployment manager.

Run the `deregisterNode` command from the `bin` directory of the administrative agent. The node must have been previously registered with the administrative agent. When you unregister a node, the node configuration is retained, but is marked as not registered with the administrative agent.

When you run the `deregisterNode` command, the command stops all running application servers on the node. You can optionally stop application servers on the node that you are unregistering prior to running the `deregisterNode` command.

If the node that you unregister had the administrative console or management Enterprise JavaBeans (EJB) applications installed prior to registering the node, they are re-enabled.

Syntax

The `deregisterNode` command syntax is as follows:

```
deregisterNode [options]
```

Parameters

The following options are available for the `deregisterNode` command:

-conntype<JSR160RMI|IPC|RMI|SOAP>

The optional connector type used to connect to the administrative agent to initiate node deregistration. The default is SOAP.

Note: You should eventually switch from the RMI connector to the JSR160RMI connector because support for the RMI connector is deprecated.

-host *host_name*

An optional parameter that specifies the host name of the administrative agent.

-nodepassword *node_password*

An optional parameter that specifies the password of the node that you are unregistering. Specify this parameter if security is on at the node and the password is different than the administrative agent password. Use this parameter with the `-nodeusername` parameter. The `-nodeusername` and `-nodepassword` parameters are used to stop all servers on the node.

-nodeusername *node_user_name*

An optional parameter that specifies the use name of the node that you are unregistering. Specify this parameter if security is on at the node and the user name is different than the administrative agent user name. Use this parameter with the `-nodepassword` parameter. The `-nodeusername` and `-nodepassword` parameters are used to stop all servers on the node.

-port *port_number*

An optional parameter that specifies the port number of the administrative agent connector port.

The default port number is 8878 for the default SOAP port of the administrative agent. SOAP is the default Java Management Extensions (JMX) connector type for the command. If you have multiple product installations or multiple profiles, the SOAP port might be different than 8878. Examine the administrative agent `SystemOut.log` file to see the current ports in use.

Note: This topic references one or more of the application server log files. Beginning in WebSphere Application Server Version 8.0 you can configure the server to use the High Performance Extensible Logging (HPEL) log and trace infrastructure instead of using `SystemOut.log`, `SystemErr.log`, `trace.log`, and `activity.log` files or native z/OS logging facilities. If you are using HPEL, you can access all of your log and trace information using the LogViewer command-line tool from your server profile bin directory. See the information about using HPEL to troubleshoot applications for more information on using HPEL.

-profilePath *profile_path*

A required parameter that specifies the path to the profile of the base node to be unregistered.

-trace

An optional parameter that provides tracing output for the `deregisterNode` command.

-profileName *profile_name*

An optional parameter that specifies the profile name of the administrative agent.

-username *username*

An optional parameter that specifies the user ID of the node to be unregistered.

-password *password*

An optional parameter that specifies the password of the node to be unregistered.

-help

An optional parameter that prints a usage statement.

-? An optional parameter that prints a usage statement.

Usage scenarios

The following examples demonstrate correct syntax. Commands are split on multiple lines for printing purposes.

```
deregisterNode.sh -conntype SOAP -port 8878
-profilePath app_server_root/profiles/AppSrv01
```

backupConfig command

The `backupConfig` command is a simple utility to back up the configuration of your node to a file.

Note:

By default, all servers on the node stop before the backup is made so that partially synchronized information is not saved. For more information about where to run this command, see *Using command line tools*. If you do not have root authority, you must specify a path for the backup file in a location where you have write permission. The backup file will be in zip format and a `.zip` extension is recommended.

In a UNIX or Linux environment, the `backupConfig` command does not save file permissions or ownership information. The `restoreConfig` command uses the current `umask` and effective user ID (EUID) to set the permissions and ownership when restoring a file. If it is required that the restored files have the original permissions and ownership, use the `tar` command (available on all UNIX or Linux systems) to back up and restore the configuration.

Note: This command uses the user ID and password information in the *profile_root/properties/*ipc.client.props file. To avoid user ID and password prompts when you use this command, add the user ID and password information to the *ipc.client.props* file.

Location

Issue the command from the *profile_root/bin* directory.

Syntax

The command syntax is as follows:

```
backupConfig.sh backup_file [options]
```

where *backup_file* specifies the file to which the backup is written. If you do not specify one, a unique name is generated.

Parameters

The following options are available for the backupConfig command:

-nostop

Tells the backupConfig command not to stop the servers before backing up the configuration

-quiet

Suppresses the progress information that the backupConfig command prints in normal mode

-logfile *file_name*

Specifies the location of the log file to which trace information is written

By default, the log file is named backupConfig.log and is created in the logs directory.

-replacelog

Replaces the log file instead of appending to the current log

-trace

Generates trace information into the log file for debugging purposes

-username *user_name*

Specifies the user name for authentication if security is enabled in the server; acts the same as the -user option

-user *user_name*

Specifies the user name for authentication if security is enabled in the server; acts the same as the -username option

-password *password*

Specifies the password for authentication if security is enabled in the server

-help

Prints a usage statement

-? Prints a usage statement

Usage

The following example creates a new file that includes the current date:

```
backupConfig.sh WebSphereConfig_2005-04-22.zip
```

The following example creates a file called myBackup.zip and does not stop any servers before beginning the backup process:

```
backupConfig.sh myBackup.zip -nostop
```

restoreConfig command

Use the restoreConfig command to restore the configuration of your node after backing up the configuration using the backupConfig command.

Note:

The restoreConfig command is a simple utility to restore the configuration of your node after backing up the configuration using the backupConfig command. By default, all servers on the node stop before the configuration restores so that a node synchronization does not occur during the restoration. If the configuration directory already exists, it is renamed before the restoration occurs. For more information about where to run this command, see Using command line tools.

If you directly make changes to the application files in the *app_server_root/installedApps* directory, a process known as "hot deployment", but do not make the same changes to the application files in the *app_server_root/config* directory, the changes might be overwritten if you use the restoreConfig command.

The backupConfig command does not save file permissions or ownership information. The restoreConfig command uses the current umask and effective user ID (EUID) to set the permissions and ownership when restoring a file. If it is required that the restored files have the original permissions and ownership, use the tar command (available on all UNIX or Linux systems) to back up and restore the configuration.

If you are using a logical directory for *app_server_root/config*, the restoreConfig command will not work.

Location

Issue the command from the *profile_root/bin* directory.

Syntax

The command syntax is as follows:

```
restoreConfig.sh backup_file [options]
```

where *backup_file* specifies the file to be restored. If you do not specify one, the command will not run.

Parameters

The following options are available for the restoreConfig command:

-help

Prints a usage statement

-location *directory_name*

Specifies the directory where the backup file is restored

The location defaults to the *app_server_root/config* directory.

-logfile *file_name*

Specifies the location of the log file to which trace information is written

By default, the log file is named *restoreConfig.log* and is created in your *logs* directory.

-nostop

Tells the restoreConfig command not to stop the servers before restoring the configuration

-password *password*

Specifies the password for authentication if security is enabled in the server

-quiet

Suppresses the progress information that the restoreConfig command prints in normal mode

-replacelog

Replaces the log file instead of appending to the current log

-trace

Generates trace information into the log file for debugging purposes

-username *user_name*

Specifies the user name for authentication if security is enabled in the server; acts the same as the -user option

-user *user_name*

Specifies the user name for authentication if security is enabled in the server; acts the same as the -username option

-? Prints a usage statement

Usage

The following example demonstrates correct syntax:

```
restoreConfig.sh WebSphereConfig_2006-04-22.zip
```

The following example restores the given file to the /tmp directory and does not stop any servers before beginning the restoration:

```
restoreConfig.sh WebSphereConfig_2006-04-22.zip -location /tmp -nostop
```

Be aware that if you restore the configuration to a directory that is different from the directory that was backed up when you performed the backupConfig command, you might need to manually update some of the paths in the configuration directory.

versionInfo command

The versionInfo command generates a report that includes a list of installed fix packs and interim fixes.

Product version information

The versionInfo tool displays important data about the product and its installed fix packs and interim fixes, such as the build version and build date. This tool is particularly useful when working with support personnel to determine the cause of any problem.

Product version reports

The following report-generation scripts display installed product information:

- versionInfo script
Lets you use parameters to create a version report.
- genVersionReport script
Generates the versionReport.html report file in the current working directory, which is usually the *app_server_root/bin* directory.

Location of the command file

Syntax for the versionInfo command

The command syntax is:

Issue the command from the `bin` directory of `app_server_root`.

Parameters

-? or /?

Displays command syntax.

-components

This parameter is deprecated and performs no action.

-componentDetail

This parameter is deprecated and performs no action.

-file *file_name*

Specifies the output file name. The report goes to standard output (stdout) by default.

-fixpacks

Adds a list of applied fixpacks to the report.

-fixpackDetail

Adds details about applied fixpacks to the report.

-format *text | html*

Selects the format of the report. The default is "text".

-help or /help

Displays command syntax.

-ifixes

Adds a list of applied ifixes to the report.

-ifixDetail

Adds details about applied ifixes to the report.

-long

Creates the long version of the report.

-maintenancePackageDetail

This option is deprecated, and it performs an action that is equivalent to `-fixpackDetail` plus `-ifixDetail`.

-maintenancePackages

This option is deprecated, and it performs an action that is equivalent to `-fixpacks` plus `-ifixes`.

-usage

Displays command syntax.

Report description

The `versionInfo` command reports the following information:

Installation information

Displays the following general information about the current installation:

- Report date and time - The date and time that the report was generated. The timestamp is formatted according to the current locale.
- Product Directory - The file path to the installation root directory defined by the `WAS_HOME` environment variable.

- Version Directory - The file path of the version directory of the current IBM WebSphere Application Server - ND installation.
- DTD Directory - The file path of the DTD directory of the current IBM WebSphere Application Server - ND installation.
- Log Directory - The file path of the log directory of the current IBM WebSphere Application Server - ND installation. The fix pack and interim fix log files are in the log directory.

Product list information

Displays a list of installed WebSphere products:

- Product ID - The product ID of the installed product.
- Status - The status of the product, either installed or uninstalled.

Installed product information

This information and the other information topic descriptions are hierarchal for each installed product, installed fix packs and interim fixes, and included APARs.

This section of the report displays the following information:

- Name - The name of the installed product.
- Version - The current version of the product. Installing or uninstalling fix packs or refresh packs modifies this version.
- ID - The product ID of the product installed, such as BASE, BASETRIAL, ND, EXPRESS, EXPRESSTRIAL, embeddedEXPRESS, IHS, XD, PLG, or CLIENT.
- Build Level - The build level of the installed product.
- Build Date - The build date of the installed product.
- Architecture - The architecture of the installed product.
- Installed Features - The features installed on the product.

Installed fix pack information: Displays the general fix pack information:

- Product ID - The ID of the product that this fix pack is for.
- Version - The version of this fix pack.
- Installation Manager Offering ID - The ID of offering or product installed using Installation Manager.
- Build Level - The build level of this fix pack.
- Build Date - The build date of this fix pack.

Installed interim fix information: Displays the general interim fix information:

- Interim Fix ID - The ID of this interim fix.
- Product ID - The ID of the product that this interim fix is for.
- Applicable Level - The level of the fix pack on which this interim fix is built.
- Installation Manager Offering ID - The ID of offering or product installed using Installation Manager.
- Build Level - The build level of this interim fix.
- Build Date - The build date of this interim fix.

Included APARs information: Displays the list of APARs fixed by this fix pack or interim fix.

Sample versionInfo report

When the WebSphere Application Server Network Deployment product has no interim fixes or fix packs applied, the versionInfo script outputs information similar to the following:

```
-----
IBM WebSphere Product Installation Status Report
-----
```

```
Report at date and time April 7, 2010 12:19:43 PM EDT
```

```

Installation
-----
Product Directory      G:\IBM\WebSphere\AppServer_v8_nd
Version Directory     G:\IBM\WebSphere\AppServer_v8_nd\properties\version
DTD Directory         G:\IBM\WebSphere\AppServer_v8_nd\properties\version\dttd
Log Directory         C:\Documents and Settings\All Users\Application Data\IBM\Installation Manager\logs

Product List
-----
ND                    installed

Installed Product
-----
Name                  IBM WebSphere Application Server - ND
Version               8.0.0.0
ID                   ND
Build Level           8.0.0.20091024_0330
Build Date            2009-10-24 03:30:00-0400
Architecture          Intel (32 bit)
Installed Features    Non-English language packages for the application server runtime environment
                    Sample applications

-----
End Installation Status Report
-----

```

genVersionReport command

The `genVersionReport` command uses the `versionInfo` command to generate the `versionReport.html` report file in the current working directory, which is usually the `bin` directory. The report includes a list of installed fix packs and interim fixes.

Product version information

The `versionInfo` tool displays important data about the product, such as the build version and build date. This tool is particularly useful when working with support personnel to determine the cause of any problem.

Product version reports

The following report-generation scripts display installed product information:

- “`versionInfo` command” on page 140
Use the `versionInfo` command to specify your own report parameters when creating a customized version report.
- `genVersionReport` command
Use the `genVersionReport` command to generate the `versionReport.html` report file in the current working directory, which is usually the `bin` directory. The report includes the list of fix packs and interim fixes.

Location of the command file

Syntax for the `genVersionReport` command

The command syntax is:

Issue the command from the `bin` directory of the `app_server_root` directory.

Report description

The `versionInfo` command reports the following information:

Installation information

Displays the following general information about the current installation:

- Report date and time - The date and time that the report was generated. The timestamp is formatted according to the current locale.
- Product Directory - The file path to the installation root directory defined by the WAS_HOME environment variable.
- Version Directory - The file path of the version directory of the current IBM WebSphere Application Server - ND installation.
- DTD Directory - The file path of the DTD directory of the current IBM WebSphere Application Server - ND installation.
- Log Directory - The file path of the log directory of the current IBM WebSphere Application Server - ND installation. The fix pack and interim fix log files are in the directory.

Product list information

Displays a list of installed WebSphere products:

- Product ID - The product ID of the installed product.
- Status - The status of the product, either installed or uninstalled.

Installed product information

This information and the other information topic descriptions are hierarchal for each installed product, installed fix packs and interim fixes, and included APARs.

This section of the report displays the following information:

- Name - The name of the installed product.
- Version - The current version of the product. Installing or uninstalling fix packs or refresh packs modifies this version.
- ID - The product ID of the product installed, such as BASE, BASETRIAL, ND, EXPRESS, EXPRESSTRIAL, embeddedEXPRESS, IHS, XD, PLG, or CLIENT.
- Build Level - The build level of the installed product.
- Build Date - The build date of the installed product.
- Architecture - The architecture of the installed product.
- Installed Features - The features installed on the product.

Installed fix pack information: Displays the general fix pack information:

- Product ID - The ID of the product that this fix pack is for.
- Version - The version of this fix pack.
- Installation Manager Offering ID - The ID of offering or product installed using Installation Manager.
- Build Level - The build level of this fix pack.
- Build Date - The build date of this fix pack.

Installed interim fix information: Displays the general interim fix information:

- Interim Fix ID - The ID of this interim fix.
- Product ID - The ID of the product that this interim fix is for.
- Applicable Level - The level of the fix pack on which this interim fix is built.
- Installation Manager Offering ID - The ID of offering or product installed using Installation Manager.
- Build Level - The build level of this interim fix.
- Build Date - The build date of this interim fix.

Included APARs information: Displays the list of APARs fixed by this fix pack or interim fix.

Sample genVersionReport report

When the WebSphere Application Server Network Deployment product has no interim fixes or fix packs applied, the genVersionReport script creates the following information in the versionReport.html report file.

On a Windows system, the report might resemble the following example:

```
-----
IBM WebSphere Product Installation Status Report
-----

Report at date and time April 7, 2010 12:19:43 PM EDT

Installation
-----
Product Directory      G:\IBM\WebSphere\AppServer_v8_nd
Version Directory     G:\IBM\WebSphere\AppServer_v8_nd\properties\version
DTD Directory         G:\IBM\WebSphere\AppServer_v8_nd\properties\version\DTD
Log Directory         C:\Documents and Settings\All Users\Application Data\IBM\Installation Manager\logs

Product List
-----
ND                    installed

Installed Product
-----
Name                  IBM WebSphere Application Server - ND
Version               8.0.0.0
ID                   ND
Build Level           8.0.0.20091024_0330
Build Date            2009-10-24 03:30:00-0400
Architecture          Intel (32 bit)
Installed Features    Non-English language packages for the application server runtime environment
                    Sample applications

-----
End Installation Status Report
-----
```

historyInfo command

The historyInfo command generates a report that includes a history of installed or uninstalled fix packs and interim fixes.

Product history information

The historyInfo tool displays important data about the product, such as the build version and build date. History information for installation and removal of fix packs and interim fixes also displays in the report. This tool is particularly useful when working with support personnel to determine the cause of any problem.

Product history reports

The following report-generation scripts display installed product information:

- historyInfo script
Lets you use parameters to create a history report.
- genHistoryReport script
Generates the historyReport.html report file in the current working directory, which is usually the bin directory.

Location of the command file

Syntax for the historyInfo command

The command syntax is:

Issue the command from the bin directory of the *app_server_root* directory.

Parameters

-? or /?

Displays command syntax.

-component *component_name*

This parameter is deprecated and performs no action.

-file *file_name*

Specifies the output file name. The report goes to standard output (stdout) by default.

-format **text** | **html**

Selects the format of the report. The default is "text".

-help or **/help**

Displays command syntax.

-maintenancePackageID *ID_of_Installation_Manager_offering*

This option is deprecated and equivalent to using **-offeringID**.

-offeringID *ID_of_Installation_Manager_offering*

Specifies the ID of the Installation Manager offering. When it is specified, the product history report displays events for only the named offering. When it is not specified, the report displays events for all offerings.

-usage

Displays command syntax.

Report description

The historyInfo command reports the following information:

Installation information

Displays the following general information about the current installation:

- Report date and time - The date and time that the report was generated. The timestamp is formatted according to the current locale.
- Product Directory - The file path to the installation root directory of the product.
- Version Directory - The file path of the version directory of the current product installation.
- DTD Directory - The file path of the DTD directory of the current installation.
- Log Directory - The file path of the log directory of the current installation. The fix pack and interim fix log files are in the directory.

Installation event information

Displays the list of installed fix packs and interim fixes as well as the following related information:

- Installation Manager Offering ID or Fix ID - The ID of the offering or product installed using Installation Manager.
- Action - The action taken.
- Version - Either the version of the interim fix or the version of the product after the action was performed.
- Log File Name - The file path of the log file generated during the event.
- Timestamp - The time when the action occurred. The time is stated in relation to GMT.
- Result - The result of the action. The result is either success, partial success, or failure.
- Installed Features - The features installed on the product.

Sample historyInfo report

The historyInfo script outputs the information similar to the following:

IBM WebSphere Product History Report

Report at date and time April 7, 2010 3:07:36 PM EDT

Installation

Product Directory G:\IBM\WebSphere\AppServer_v8_nd
Version Directory G:\IBM\WebSphere\AppServer_v8_nd\properties\version
DTD Directory G:\IBM\WebSphere\AppServer_v8_nd\properties\version\dttd
Log Directory C:\Documents and Settings\All Users\Application Data\IBM\Installation Manager\logs

Installation Event

Install Manager Offering ID com.ibm.websphere.ND.v80
Action install
Version 8.0.0.0
Log File Name C:\Documents and Settings\All Users\Application Data\IBM\Installation Manager\logs\20091026_1527.xml
Timestamp 2009-10-26 17:35:27-0400
Result success
Installed Features Sample applications

Installation Event

Install Manager Offering ID com.ibm.websphere.ND.v80
Action install
Version 8.0.0.0
Log File Name C:\Documents and Settings\All Users\Application Data\IBM\Installation Manager\logs\20091027_1347.xml
Timestamp 2009-10-27 13:48:03-0400
Result success
Installed Features Sample applications
Non-English language packages for the application server runtime environment

Installation Event

Install Manager Offering ID com.ibm.websphere.ND.v80
Action update
Version 8.0.0.7
Log File Name C:\Documents and Settings\All Users\Application Data\IBM\Installation Manager\logs\20091027_1347.xml
Timestamp 2009-10-27 13:53:48-0400
Result success
Installed Features Non-English language packages for the application server runtime environment
Sample applications

Installation Event

Fix ID com.ibm.websphere.ND.v80.FP8007.fix.1
Action install
Version 8.0.0.7
Log File Name C:\Documents and Settings\All Users\Application Data\IBM\Installation Manager\logs\20091027_1347.xml
Timestamp 2009-10-27 15:09:43-0400
Result success

Installation Event

Install Manager Offering ID com.ibm.websphere.ND.v80
Action uninstall
Version 8.0.0.7
Log File Name C:\Documents and Settings\All Users\Application Data\IBM\Installation Manager\logs\20091027_1347.xml
Timestamp 2009-10-27 16:01:52-0400
Result success
Installed Features Non-English language packages for the application server runtime environment

End History Report

genHistoryReport command

The genHistoryReport command generates the historyReport.html report file in the current working directory, which is usually the bin directory. The report includes a list of installed or uninstalled fix packs and interim fixes. The genHistoryReport script invokes the historyInfo script specifying the correct parameters to place the information generated into an HTML file in the current directory.

Product history information

The historyInfo tool displays historical data about the product and the installation and removal of fix packs and interim fixes for the product. This tool is particularly useful when working with support personnel to determine the cause of any problem.

Product history reports

The following report-generation scripts display installed product information:

- “historyInfo command” on page 145
Lets you use parameters to create a history report.
- genHistoryReport script
Generates the `historyReport.html` report file in the current working directory, which is usually the `bin` directory. The report includes a list of fix packs and interim fixes.

Location of the command file

Syntax for the genHistoryReport command

The command syntax is:

Issue the command from the `bin` directory of the `app_server_root` directory.

Report description

The `historyInfo` command generates the report. The `genHistoryReport` command calls the `historyInfo` command with a set of report parameters that reports the following information:

Installation information

Installation information displays the following general information about the current installation:

- Report date and time - The date and time that the report was generated. The timestamp is formatted according to the current locale.
- Product Directory - The file path to the installation root directory of the product.
- Version Directory - The file path of the `version` directory of the current product installation.
- DTD Directory - The file path of the `DTD` directory of the current installation.
- Log Directory - The file path of the `log` directory of the current installation. The fix pack and interim fix log files are in the directory.

Installation event information

Installation event information displays the list of installed fix packs and interim fixes as well as the following related information:

- Installation Manager Offering ID or Fix ID - The ID of the offering or product installed using Installation Manager.
- Action - The action taken.
- Version - Either the version of the interim fix or the version of the product after the action was performed.
- Log File Name - The file path of the log file generated during the event.
- Timestamp - The time when the action occurred. The time is stated in relation to GMT.
- Result - The result of the action. The result is either success, partial success, or failure.
- Installed Features - The features installed on the product.

Sample historyInfo report

The `genHistoryReport` script creates the following information in the `historyReport.html` report file:

```
-----  
IBM WebSphere Product History Report  
-----
```

```
Report at date and time April 7, 2010 3:07:36 PM EDT
```

```
Installation  
-----
```



```

Product Directory   G:\IBM\WebSphere\AppServer_v8_nd
Version Directory  G:\IBM\WebSphere\AppServer_v8_nd\properties\version
DTD Directory       G:\IBM\WebSphere\AppServer_v8_nd\properties\version\DTD
Log Directory       C:\Documents and Settings\All Users\Application Data\IBM\Installation Manager\logs

```

Installation Event

```

-----
Install Manager Offering ID  com.ibm.websphere.ND.v80
Action                       install
Version                      8.0.0.0
Log File Name                 C:\Documents and Settings\All Users\Application Data\IBM\Installation Manager\logs\20091026_1527.xml
Timestamp                    2009-10-26 17:35:27-0400
Result                        success
Installed Features            Sample applications

```

Installation Event

```

-----
Install Manager Offering ID  com.ibm.websphere.ND.v80
Action                       install
Version                      8.0.0.0
Log File Name                 C:\Documents and Settings\All Users\Application Data\IBM\Installation Manager\logs\20091027_1347.xml
Timestamp                    2009-10-27 13:48:03-0400
Result                        success
Installed Features            Sample applications
                               Non-English language packages for the application server runtime environment

```

Installation Event

```

-----
Install Manager Offering ID  com.ibm.websphere.ND.v80
Action                       update
Version                      8.0.0.7
Log File Name                 C:\Documents and Settings\All Users\Application Data\IBM\Installation Manager\logs\20091027_1347.xml
Timestamp                    2009-10-27 13:53:48-0400
Result                        success
Installed Features            Non-English language packages for the application server runtime environment
                               Sample applications

```

Installation Event

```

-----
Fix ID                      com.ibm.websphere.ND.v80.FP8007.fix.1
Action                       install
Version                      8.0.0.7
Log File Name                 C:\Documents and Settings\All Users\Application Data\IBM\Installation Manager\logs\20091027_1347.xml
Timestamp                    2009-10-27 15:09:43-0400
Result                        success

```

Installation Event

```

-----
Install Manager Offering ID  com.ibm.websphere.ND.v80
Action                       uninstall
Version                      8.0.0.7
Log File Name                 C:\Documents and Settings\All Users\Application Data\IBM\Installation Manager\logs\20091027_1347.xml
Timestamp                    2009-10-27 16:01:52-0400
Result                        success
Installed Features            Non-English language packages for the application server runtime environment

```

End History Report

managesdk command

The managesdk command provides the names of software development kits that are used by the product.

Note: Use the managesdk command to:

- List the software development kit (SDK) names that are available to a product installation.
- List the SDK names that a specified profile is currently configured to use.
- For each profile in a product installation, list the SDK names that the profile is currently configured to use.
- Enable a profile to use a specified SDK name.
- Enable all profiles in an installation to use a specified SDK name.
- Get the SDK name that is used to configure new profiles.
- Change the default SDK name that profiles use.
- Get the SDK name that is used by scripts called from a product bin directory.
- Change the SDK name that scripts in a product bin directory use by default. The SDK name is used when no existing profile name is specified and the default profile name is not applicable.

The command file is located in the *app_server_root/bin* directory, *app_client_root/bin* directory, and *plugins_root/bin* directory.

Attention: If the `managesdk` command is used to change the SDK for a profile from a 31-bit (z/OS) or 32-bit (IBM i) SDK to a 64-bit SDK, and you are using third-party resource adapters, consider the following information to avoid potential problems. This information does not apply to any of the built-in resource adapters shipped with the WebSphere Application Server product, including the IBM WebSphere Relational Resource Adapter, the IBM WebSphere MQ Resource Adapter, or the IBM SIB JMS Resource Adapter as they have been fully tested to work with all IBM SDKs. Because resource adapters can use non-Java libraries containing platform-specific native code, it is possible that changing the SDK from 31-bit (z/OS) or 32-bit (IBM i) to 64-bit, or from 64-bit to 31-bit or 32-bit, might result in the resource adapter not functioning properly. If a third-party resource adapter is installed, either stand-alone or embedded in an enterprise application, on a server for which you intend to change the SDK, verify with the supplier of that resource adapter that any native libraries it uses are compatible with the selected SDK.

Syntax

Use the following command syntax with the `managesdk` command:

```
managesdk -task [-parameter] [value]
```

The command-line tool validates that the requested task contains the required parameters and values. Parameters are not case-sensitive. However, values are case-sensitive. You must type values with the correct capitalization because the command-line tool does not validate the capitalization of the parameter values. Incorrect results can occur when the parameter value is not typed correctly.

Parameters

The following `-task` options are available for the `managesdk` command:

-help

Displays detailed information about the parameters or values of each `managesdk` task. The following example uses the `help` parameter with the `managesdk` command:

```
app_server_root/bin/managesdk.sh -help
```

The output from the `help` option describes the required and optional parameters.

-listAvailable [-verbose]

Displays a list of all SDK names available to the product installation. When the `-verbose` option is also specified, a list of properties for each SDK name also is displayed. The following example uses the `-listAvailable -verbose` parameters with the `managesdk` command:

```
app_server_root/bin/managesdk.sh -listAvailable -verbose
```

The output is a list of all SDK names that the product installation can use, along with a list of the properties associated with each SDK name.

-listEnabledProfile [-profileName profile_name] [-verbose]

Displays a list of all SDK names that a specified profile, and its node and servers, is currently configured to use. When the `-verbose` option is also specified, a list of properties for each SDK name also is displayed. The following example uses the `-listEnabledProfile -profileName` and `-verbose` parameters with the `managesdk` command:

```
app_server_root/bin/managesdk.sh -listEnabledProfile -profileName AppSrv02 -verbose
```

The output is a list of all SDK names that the specified profile can use, along with a list of the properties associated with each SDK name.

-listEnabledProfileAll [-verbose]

Displays a list of all profiles in an installation and the SDK names that each profile, and its node and servers, is currently configured to use. When the `-verbose` option is also specified, a list of properties

for each SDK name also is displayed. The following example uses the `-listEnabledProfileAll` and `-verbose` parameters with the `managesdk` command:

```
app_server_root/bin/managesdk.sh -listEnabledProfileAll -verbose
```

The output is a list of all profiles in a product installation with all SDK names that each profile can use, along with a list of the properties associated with each SDK name.

-enableProfile [-profileName profile_name] [-sdkname sdkName] [-enableServers] [-user user_name] [-password password_value]

Enables a profile to use a specified SDK name. The `-profileName` parameter specifies the profile and the `-sdkname` parameter specifies the SDK name. The command enables the profile and the node-level default SDK of the profile to use the specified SDK name. Unless the `-enableServers` option is used, the command does not change server-level SDK settings. If the `-enableServers` option is used, all server-level SDK settings are cleared, enabling all servers to use the node-level default SDK.

The following conditions apply when the `managesdk` command is run:

- If the profile is a federated node or a deployment manager node, the deployment manager must be running when the `managesdk` command attempts to update the profile. When enabling the SDK for a node, run the `managesdk` command from the `/bin` directory of the product installation to which the node belongs or from the `/bin` directory of the profile that contains the node you want to update.
- A connection to the deployment manager must exist using a supported connector protocol in the following order of preference:
 1. SOAP
 2. Inter-Process Communications (IPC)
 3. Remote Method Invocation (RMI)

If the SOAP protocol is enabled, the `managesdk` command uses the SOAP protocol. If the SOAP protocol is not enabled but the IPC protocol is enabled, the command uses the IPC protocol. If neither the SOAP nor the IPC protocol are enabled, then the command uses the RMI protocol.

- You must provide the administrative user name and password with the `managesdk` command for each profile that contains a federated node or deployment manager node in a cell with security enabled. If you do not specify the `-user` and `-password` parameters, the `managesdk` command might fail or stop processing. The topic on configuring security with scripting provides connector protocol specific instructions on how to save user name and password values.
- When enabling the SDK for a deployment manager, only the deployment manager server is enabled. None of the managed nodes of the deployment manager are enabled to use the specific SDK.

The following example uses `-enableProfile`, `-profileName`, `-sdkname`, and `-enableServers` with the `managesdk` command:

```
app_server_root/bin/managesdk.sh -enableProfile -profileName AppSrv02 -sdkname 1.6_32 -enableServers
app_server_root/bin/managesdk.sh -enableProfile -profileName AppSrv02 -sdkname 1.6_64 -enableServers
```

The output is a message that indicates whether the specified profile was successfully updated and is now enabled to use the specified SDK, or whether problems were encountered that prevented the profile from being successfully updated.

-enableProfileAll [-sdkname sdkName] [-enableServers] [-user user_name] [-password password_value]

Enables all profiles in an installation to use a specified SDK name. The `-sdkname` parameter specifies the SDK name. The command enables all profiles and the node-level default SDK of each profile to use the specified SDK name. Unless the `-enableServers` option is used, the command does not change server-level SDK settings. If the `-enableServers` option is used, all server-level SDK settings are cleared, enabling all servers to use the node-level default SDK.

The following conditions apply when the `managesdk` command is run:

- If the profile is a federated node or a deployment manager node, the deployment manager must be running when the `managesdk` command attempts to update the profile. When enabling the SDK for a node, run the `managesdk` command from the `/bin` directory of the product installation to which the node belongs or from the `/bin` directory of the profile that contains the node you want to update.
- A connection to the deployment manager must exist using a supported connector protocol in the following order of preference:
 1. SOAP
 2. Inter-Process Communications (IPC)
 3. Remote Method Invocation (RMI)

If the SOAP protocol is enabled, the `managesdk` command uses the SOAP protocol. If the SOAP protocol is not enabled but the IPC protocol is enabled, the command uses the IPC protocol. If neither the SOAP nor the IPC protocol are enabled, then the command uses the RMI protocol.

- You must provide the administrative user name and password with the `managesdk` command for each profile that contains a federated node or deployment manager node in a cell with security enabled. If you do not specify the `-user` and `-password` parameters, the `managesdk` command might fail or stop processing. The topic on configuring security with scripting provides connector protocol specific instructions on how to save user name and password values.

Note: Do not use the `-enableProfileAll` option unless automatic prompting is disabled for SOAP, IPC, and RMI connections to the deployment managers of cells that have any federated node or deployment manager with security enabled. Automatic prompting causes the `managesdk` command to fail or stop processing.

- When enabling the SDK for a deployment manager, only the deployment manager server is enabled. None of the managed nodes of the deployment manager are enabled to use the specific SDK.

The following example uses `-enableProfileAll`, `-sdkname`, and `-enableServers` with the `managesdk` command:

```
app_server_root/bin/managesdk.sh -enableProfileAll -sdkname 1.6_64 -enableServers
```

The output is a message for each profile that indicates whether the profile was successfully updated and is now enabled to use the specified SDK, or whether problems were encountered that prevented the profile from being successfully updated.

-getNewProfileDefault [-verbose]

Displays the SDK name that is currently configured for all profiles that are created with the `manageprofiles` command. When the `-verbose` option is also specified, properties information for the single SDK name also is displayed.

The following example uses the `-getNewProfileDefault -verbose` parameters with the `managesdk` command:

```
app_server_root/bin/managesdk.sh -getNewProfileDefault -verbose
```

After the command runs, the new profile default SDK name is displayed.

-setNewProfileDefault [-sdkname sdkName]

Changes the SDK name that is currently configured for all profiles that are created with the `manageprofiles` command. The `-sdkname` parameter specifies the default SDK name to use. The `sdkName` value must be an SDK name that is enabled for the product installation.

The following example uses the `-setNewProfileDefault -sdkname` parameters with the `managesdk` command:

```
app_server_root/bin/managesdk.sh -setNewProfileDefault -sdkname 1.6_64
```

After the command runs, the new profile default SDK name is displayed.

-getCommandDefault [-verbose]

Displays the SDK name that script commands in the `app_server_root/bin`, `app_client_root/bin`, or

plugins_root/bin directory are enabled to use when no existing profile name is specified or when the default profile name is used. When the `-verbose` option is also specified, properties information for the single SDK name also is displayed.

The following example uses the `-getCommandDefault -verbose` parameters with the `managesdk` command:

```
app_server_root/bin/managesdk.sh -getCommandDefault -verbose
```

-setCommandDefault [-sdkname sdkName]

Changes the SDK name that script commands in the *app_server_root/bin*, *app_client_root/bin*, or *plugins_root/bin* directory are enabled to use when no existing profile name is specified or when the default profile name is used. The `-sdkname` parameter specifies the SDK name to use for commands. The *sdkName* value must be an SDK name that is enabled for the product installation.

The following example uses the `-getCommandDefault -sdkname` parameters with the `managesdk` command:

```
app_server_root/bin/managesdk.sh -setCommandDefault -sdkname 1.6_64
```

The following special parameter options are available with task parameters of the `managesdk` command:

-debug

Use this option with any `-task` parameter to enable additional debugging information in the command output.

-quiet

Use this option with any `-task` parameter to suppress most messages in the command output.

-sdkname

Use this option with a `-set` task parameter to specify an SDK name that is enabled for the product installation; for example:

- `-sdkname 1.6_31`

-verbose

Use this option with any `-list` or `-get` task parameter to provide additional information, such as SDK properties in the command output.

GenPluginCfg command

The `GenPluginCfg` command is used to regenerate the WebSphere web server plug-in configuration file, `plugin-cfg.xml`.

For more information about where to run this command, see the [Using command tools](#) article.

CAUTION:

Regenerating the plug-in configuration can overwrite manual configuration changes that you might want to preserve. Before performing this task, understand its implications as described in the [Communicating with web servers](#) topic in the *Setting up the application serving environment* PDF.

Note: You must delete the `plugin-cfg.xml` file in the *profile_root/config/cells* directory before you use this command. Otherwise, configuration changes do not persist to the `plugin-cfg.xml` file.

Note: You can update the *global* `plugin-cfg.xml` file using the administrative console or running the `GenPluginCfg` command for all of the clusters in a cell. However, you must delete the `config/cells/plugin-cfg.xml` file before you update the *global* `plugin-cfg.xml` file. If you do not delete the `config/cells/plugin-cfg.xml` file, only the new properties and their values are added to the *global* `plugin-cfg.xml` file. Any updates to existing plug-in property values are not added to the *global* `plugin-cfg.xml` file.

Syntax

To regenerate the plug-in configuration perform one of the following:

- Issue the following command:

```
app_server_root/bin/GenPluginCfg.sh
```

This method for regenerating the plug-in configuration creates a `plugin-cfg.xml` file in UTF-8 format, which is the proper format for execution in a z/OS environment.

When the `GenPluginCfg` command is issued with the option `-webserver.name webserverName`, `wsadmin` generates a plug-in configuration file for the web server. The settings in the generated configuration file are based on the list of applications that are deployed on the web server. When this command is issued without the option `-webserver.name webserverName`, the plug-in configuration file is generated based on topology.

Parameters

The following options are available for the `GenPluginCfg` command:

-config.root configroot_dir

Defaults to `CONFIG_ROOT`. The `setupCmdLine` command is invoked to get this environment variable.

-cell.name cell

Defaults to `WAS_CELL`. The `setupCmdLine` command is invoked to get this environment variable.

-node.name node

Defaults to `WAS_NODE`. The `setupCmdLine` command is invoked to get this environment variable.

-webserver.name webserver1

Required for creating plug-in configuration file for a given Web server.

-propagate yes/no

Applicable only when the `webserver.name` option is specified and the web server is local. Otherwise, you must manually copy the `plugin-cfg.xml` file from `app_server_root/profiles/profile_name/config/cells/cell_name/nodes/node_name/servers/web_server_name` to `plugins_root/config/web_server_name` in the remote web server plugins directory. The default value is `no`.

-propagateKeyring yes/no

Applicable only when the option `webserver.name` is specified and the web server is local. Defaults to `no`.

-cluster.name cluster1,cluster2 | ALL

Optional list of clusters. Ignored when the option `webserver.name` is specified.

-server.name server1,server2

Optional list of servers. Required for single server plug-in generation. Ignored when the option `webserver.name` is specified.

-output.file.name file_name

Defaults to the `configroot_dir/plugin-cfg.xml` file. Ignored when the option `webserver.name` is specified.

-destination.root root

Installation root of the machine configuration is used on. Ignored when the option `webserver.name` is specified.

-destination.operating.system windows/unix

Operating system of the machine configuration is used on. Ignored when the option `webserver.name` is specified.

- force yes**
Creates a new configuration instead of attempting to merge with an existing configuration when command is issued for a cell-wide file generation.
- debug yes/no**
Defaults to no.
- help**
Prints a usage statement.
- ?** Prints a usage statement.

Usage scenario

To generate a plug-in configuration for all of the clusters in a cell:

```
GenPluginCfg.sh -cell.name NetworkDeploymentCell
```

To generate a plug-in configuration for a single server:

```
GenPluginCfg.sh -cell.name BaseApplicationServerCell -node.name appServerNode -server.name appServerName
```

To generate a plug-in configuration file for a web server:

```
GenPluginCfg.sh -cell.name BaseApplicationServerCell -node.name webserverNode -webserver.name webserverName
```

EARExpander command

Use the **EARExpander** command to expand an enterprise archive file (EAR) into a directory to run the application in that EAR file.

You can collapse a directory containing application files into a single EAR file. You can type **EARExpander** with no arguments to learn more about its options. For more information about where to run this command, see the topic on using command tools.

Restriction: Do not include a pound sign (#) in the name of files that are packaged within an application archive. Due to internal processing, the application server fails to correctly deploy the application when a pound sign is included in a file name within the application archive. When this failure occurs, an exception might occur when the application is being processed. Also, parts of the application might be missing after the application is deployed. To address this issue, rename any file names within the application archive so that they do not contain a pound sign.

Syntax

The command syntax is as follows:

```
EarExpander -ear earName -operationDir dirName -operation
<expand | collapse> [-expansionFlags <all|war>]
```

Parameters

The following options are available for the **EARExpander** command:

- ear**
Specifies the name of the input EAR file for the expand operation or the name of the output EAR file for the collapse operation.
- operationDir**
Specifies the directory where the EAR file is expanded or specifies the directory from where files are collapsed.

-operation <expand | collapse>

The expand value expands an EAR file into a directory structure required by the WebSphere Application Server run time. The collapse value creates an EAR file from an expanded directory structure.

-expansionFlags <all | war>

(Optional) The all value expands all files from all of the modules. The war value only expands the files from Web archive file (WAR) modules.

Usage scenario

The following examples demonstrate correct syntax:

```
EARExpander.sh -ear /WebSphere/AppServer/installableApps/DefaultApplication.ear
-operationDir /MyApps -operation expand -expansionFlags war
```

```
EARExpander.sh -ear /backup/DefaultApplication.ear
-operationDir /MyAppsDefaultApplication.ear -operation collapse
```

Return codes

The EARExpander command has the following return codes.

Table 11. Return codes and their descriptions. The return code indicates the success of the operation.

Return code	Description
-1	A syntax error exists.
0	The command ran successfully.
1	An error occurred.
2	An exception occurred.

revokeCertificate command

The revokeCertificate command uses an implementation class that is passed to communicate with a certificate authority (CA) server to revoke a certificate. Processing this command sends a revocation request to the CA server to mark this certificate as revoked.

Location

Issue the command from the *profile_root/bin* directory.

Syntax

The command syntax is as follows:

(The command is split on multiple lines for printing purposes.)

```
revokeCertificate.sh -host<caHost> -port<caPort> -username<caUserName> -password<caPassword>
-revocationPassword<revocationPassword> -keystoreAlias<keystoreAlias> -alias<certificateAlias>
-pkiImplClass<customCAClient>[options]
```

Required Parameters

The following required parameter are used with the revokeCertificate command:

-host caHost

Specifies the target certificate authority host to which the request is sent.

-port caPort

Specifies the target port to connect to.

-username *caUserName*

Specifies the user name used to gain access to the certificate authority.

-password *caPassword*

Specifies the password used to authenticate with the certificate authority.

-revocationPassword *revocationPassword*

Specifies the password that is to be set on the certificate returned by the certificate authority. The revocation password is sent to the certificate authority during each request and is associated with each certificate that is issued. To later revoke a certificate, the same revocation password must be sent during a revokeCertificate request.

keyStoreAlias *keyStoreAlias*

Specifies the name of the keystore that is located in the ssl.client.props file for the profile to which the CA signed certificate is added. This file is usually the ClientDefaultKeyStore file for either a managed or unmanaged environment.

-alias *certificateAlias*

Specifies The alias of the certificate request to be revoked. The certificate is stored in the keystore specified on the request.

-pkImplClass *custom CA Client*

A class that implements the WSPKIClient interface. The implementation class handles all the communication to the CA server. This can be a custom class or a class provided with the product.

Optional Parameters

The following options are available for the revokeCertificate command:

-revocationReasonUsage *revocation reason*

The reason for revoking the certificate. The default value is “unspecified”.

-customAttrs *customAttr1=value;customAttr2=value;...*

A semi-colon separated list of custom name=value pairs to be passed in to the custom implementation class. This parameter provides a way to pass custom information to the implementation class. The ‘attr’ and ‘value’ pairs are converted to a hash map and passed to the implementation class.

-logfile *filename*

Overrides the default trace file. By default, the trace appears in the profiles/profile_name/log/caClient.log. file.

-trace

When specified, **-trace** enables tracing of the trace specification necessary to debug this component. By default, the trace appears in the profiles/profile_name/log/caClient.log file.

-replaceLog

An option to cause the existing trace file to be replaced when the command is executed. -quit

-quiet

An option to suppress most messages from printing out on the console.

-help

The option to print a usage statement

-? The option to print a usage statement

Usage

The following example performs a revokeCertificate:

```
revokeCertificate.sh -host localhost -port 1077
-username pkuser -password webspherepki -alias cert1 -keyStoreAlias ClientDefau
ltKeyStore -revocationPassword webspherepki
CWPKI0403I: Trace is being logged to the following location:
```

```
C:\opt\WebSphere\AppClient\logs\caClient.log
CWPKI0461I: Revoking a CA signed certificate.
CWPKI0462I: CA Signed Certificate Revoked [Issued By: O=IBM, C=US, Issued To:
CN=mycn, O=ibm, C=us, Not Before: Thu Feb 22 09:07:53 CST 2007, Not
After: Sat Feb 16 10:09:19 CST 2008] for reason: unspecified
```

requestCertificate command

The requestCertificate command uses an implementation class that is passed in to communicate with a certificate authority (CA) server to request a CA signed certificate. The command then adds the certificate to a supplied keystore.

The requestCertificate command can use a predefined certificate request that was created with the createCertRequest command or it creates the certificate request itself. Depending on the CA server that the command is targeted for, a completed signed request can be returned; or the CA server could accept the request and require that a call be made at a later time to get the certificate with the queryCertificate command.

Location

Issue the command from the *profile_root/bin* directory.

Syntax

The command syntax is as follows:

(The following command is split on multiple lines for printing purposes.)

```
requestCertificate.sh -host<caHost> -port<caPort> -username<caUserName> -password<caPassword>
-revocationPassword<revocationPassword> -keystoreAlias<keystoreAlias>
-pkiImplClass<customCAClient>[options]
```

Required Parameters

The following required parameters are used with the requestCertificate command:

-host *caHost*

Specifies the target certificate authority host to which the request will be sent.

-port *caPort*

Specifies the target port on which to connect.

-username *caUserName*

The user name used to gain access to the certificate authority.

-password *caPassword*

The password used to authenticate with the certificate authority.

-revocationPassword *revocationPassword*

The password that is to be set on the certificate returned by the certificate authority. The revocation password is sent to the certificate authority during each request and is associated with each certificate that is issued. To later revoke a certificate, the same revocation password must be sent during a revokeCertificate request.

keyStoreAlias *keyStoreAlias*

The name of the keystore that is located in the ssl.client.props file for the profile to which the CA signed certificate is added. This will typically be the ClientDefaultKeyStore file for either a managed or unmanaged environment.

-pkiImplClass *custom CA client*

A class that implements the WSPKIClient interface. The implementation class handles all the communication to the CA server. This could be a custom class or a class provided with the product.

Optional Parameters

The following options are available for the requestCertificate command:

-certReqPath *certificate request file*

A path to an existing PKCS10 certificate request saved in a BASE64 encoded file. If no request is specified a PKCS10 certificate request will be created automatically. In that case it is required to specify a “subjectDN” and “alias” option. By default the request will be created in the same location as the keyStore specified in the request. This will typically be in the /profile_name/etc/ directory for either a managed or unmanaged environment.

-subjectDN *subjectDN*

The distinguished name to be used for the PKCS10 certificate request. The distinguished name must contain the CN field. This option is only required if you do not specify the –certReqPath option, or if the –certReqPath option points to a file that does not exist.

-alias *certificateAlias*

The alias used to store the PKCS10 certificate request certificate in the keyStore specified on the request. Note that the CA signed certificate is stored under the same alias and will replace the cert request certificate when received. This option is only required if you do not specify the –certReqPath option, or if the –certReqPath option points to a file that does not exist.

-keySize *key size*

The size of the key. This option is only used valid if creating a PKCS10 certificate request in-band. Default size is 1024. Valid values include 512, 1024, and 2048

-keyUsage

A semi-colon separated list of extended key usage strings. This option is only valid if creating a PKCS10 certificate request in-band.

-extKeyUsage *extKeyUse1;extKeyUse2;...*

A semi-colon separated list of extended key usage strings. This option is only valid if creating a PKCS10 certificate request in-band.

-customAttrs *customAttr1=value;customAttr2=value;...*

A semi-colon separated list of custom name=value pairs to be passed in to the custom implementation class. This provides a way to pass custom information to the implementation class. The ‘attr’ and ‘value’ pairs will be converted to a hash map and passed along to the implementation class.

-retryInterval *retry interval*

The time period in seconds between retries of queries to the CA for a CA signed certificate.

-retryLimit *retry limit*

The total number of times to retry a query request to the CA.

-logfile *filename*

Overrides the default trace file. By default, the trace appears in the profiles/profile_name/log/caClient.log. file.

-trace

When specified, this enables tracing of the trace specification necessary to debug this component. By default, the trace will appear in the profiles/profile_name/log/caClient.log file.

-replaceLog

Causes the existing trace file to be replaced when the command is executed. -quit

-quiet

Suppresses most messages from printing out on the console.

- help
Prints a usage statement
- ? Prints a usage statement

Usage

The following example performs a requestCertificate:

```
requestCertificate.sh -host localhost -port 1077
  -username pkiuser -password webspherepki -revocationPassword webspherepki -keyS
toreAlias ClientDefaultKeyStore -certReqPath C:\opt\WebS
phere\AppClient\etc\certReq26924.req -trace
CWPKI0403I: Trace is being logged to the following location:
           C:\opt\WebSphere\AppClient\logs\caClient.log
CWPKI0455I: Requesting a CA signed certificate.
CWPKI0456I: CA Signed Certificate Received [Issued By: O=IBM, C=US, Issued To:
           CN=mycn, O=ibm, C=us, Not Before: Thu Feb 22 09:07:53 CST 2007, Not
           After: Sat Feb 16 10:09:19 CST 2008]
```

createCertRequest command

The createCertRequest command creates a PKCS10 certificate request and stores it in a client keystore so that it can be used to send to a certificate authority (CA) server using the requestCertificate command line utility.

Location

Issue the command from the *profile_root/bin* directory.

Syntax

The command syntax is as follows:

```
createCertRequest.sh -keyStoreAlias<keystoreAlias> -subjectDN<subjectDN> -alias<certificateAlias> [options]
```

Required Parameters

The following required parameter are used with the createCertRequest command:

-keyStoreAlias *keyStoreAlias*

Specifies the name of the keystore that is located in the ssl.client.props file for the profile to which the CA signed certificate is added. This is the name of the ClientDefaultKeyStore file for either a managed or unmanaged environment.

-subjectDN *subjectDN*

Specifies the distinguished name (DN) to be used for the PKCS10 certificate request. The DN must contain the CN, O and C fields at a minimum.

-alias *certificateAlias*

Specifies the alias used to store the PKCS10 certificate request certificate in the keystore specified on the request.

Note: the CA signed certificate is stored under the same alias and replaces the cert request certificate when received.

Optional Parameters

The following options are available for the createCertRequest command:

-keySize *key size*

An option that specifies the size of the key. This option is only used valid if creating a PKCS10 certificate request in-band. Valid values include 512, 1024 2048, 4096 and 8192. Thd default size is 2048.

-certValidity *valid days*

The time period of certificate validity. Time period is measured from current date. This option is only valid if creating a PKCS10 certificate request in-band. Default value is 365 days.

-subjectAltNames *altName1;altName2;...*

A semi-colon separated list of subject alternate names. This option is only used if creating a PKCS10 certificate request in-band.

-keyUsage *keyUse1;keyUse2;...*

A semi-colon separated list of key usage strings. This option is only valid if creating a PKCS10 certificate request in-band.

-extKeyUsage *extKeyUse1;extKeyUse2;...*

A semi-colon separated list of extended key usage strings. This option is only valid if creating a PKCS10 certificate request in-band.

-logfile *filename*

The logfile that overrides the default trace file. By default, the trace appears in the profiles/profile_name/log/caClient.log. file.

-trace

When specified, **-trace** enables tracing of the trace specification necessary to debug this component. By default, the trace will appear in the profiles/profile_name/log/caClient.log file.

-replaceLog

An option to cause the existing trace file to be replaced when the command is executed.

-quiet

An option to suppress most messages from printing out on the console.

-help

The option to print a usage statement

-? The option to print a usage statement

Usage

The following example creates a PKCS10 certificate request for a client that can be used to send to a CA :

```
createCertRequest.sh -keyStoreAlias ClientDefaultKeyStore -subjectDN CN=mycn,o=ibm,c=us -alias cert1
CWPKI0403I: Trace is being logged to the following location:
           C:\opt\WebSphere\AppClient\logs\caClient.log
CWPKI0422I: Generating a PKCS10 certificate request
CWPKI0421I: A PKCS10 certificate was successfully created. The request
           is stored in file:
           C:\opt\WebSphere\AppClient\etc\certReq26924.req
```

queryCertificate command

The queryCertificate command uses an implementation class that is passed to communicate with a certificate authority (CA) server and query a certificate.

The queryCertificate command checks to see if the certificate is complete. If the certificate is complete, then the CA certificate is stored in the client keystore. If the certificate is not complete, the certificate request remains in the keystore and the queryCertificate command can be called at some later time to determine if the certificate is complete.

Location

Issue the command from the *profile_root/bin* directory.

Syntax

The command syntax is as follows:

(The command is split on multiple lines for printing purposes.)

```
queryCertificate.sh -host<caHost> -port<caPort> -username<caUserName> -password<caPassword>
-alias<certificateAlias> -keystoreAlias<keystoreAlias>
-pkiImplClass<customCAClient>[options]
```

Required Parameters

The following required parameter are used with the queryCertificate command:

-host *caHost*

Specifies the target certificate authority host to which the request is sent.

-port *caPort*

Specifies the target port to connect to.

-username *caUserName*

Specifies the user name used to gain access to the certificate authority.

-password *caPassword*

Specifies the password used to authenticate with the certificate authority.

-alias *certificateAlias*

Specifies The alias of the certificate to be queried.

keyStoreAlias*keyStoreAlias*

Specifies the name of the keystore that is located in the *ssl.client.props* file for the profile to which the CA signed certificate is added. This name is the *ClientDefaultKeyStore* file for either a managed or unmanaged environment.

-pkiImplClass *custom CA client*

A class that implements the *WSPKIClient* interface. The implementation class handles all the communication to the CA server. This can be a custom class or a class provided with the product.

Optional Parameters

The following options are available for the queryCertificate command:

-customAttrs *customAttr1=value;customAttr2=value;...*

A semi-colon separated list of custom name=value pairs to be passed in to the custom implementation class. This parameter provides a way to pass custom information to the implementation class. The 'attr' and 'value' pairs are converted to a hash map and passed along to the implementation class.

-retryInterval *retry interval*

The time period in seconds between retries of queries to the CA server for a CA signed certificate.

-retryLimit *retry limit*

The total number of times to retry a query request to the CA server.

-logfile *filename*

The logfile that overrides the default trace file. By default, the trace appears in the *profiles/profile_name/log/caClient.log* file.

-trace

When specified, **-trace** enables tracing of the trace specification necessary to debug this component. By default, the trace appears in the profiles/profile_name/log/caClient.log file.

-replaceLog

An option to cause the existing trace file to be replaced when the command is executed.

-quiet

An option to suppress most messages from printing out on the console.

-help

The option to print a usage statement

Usage

The following example performs a queryCertificate:

```
queryCertificate.sh -host localhost -port 1077 -  
username pkiuser -password webspherepki -alias C:\opt\WebSphere\AppClient\  
etc\certReq26924.req -keyStoreAlias ClientDefaultKeyStore  
CWPKI0403I: Trace is being logged to the following location:  
C:\opt\WebSphere\AppClient\logs\caClient.log  
CWPKI0418E: The following error occurred while querying the CA for a signed  
certificate: CWPKI0463I: Action "query" not supported by this  
implementation.
```

restoreJobManager command

The restoreJobManager command restores the job manager database from a previous backup.

To perform the restore, the following situations must exist:

- The job manager server must be stopped.
- No other process can have a lock on the job manager database.

For more information about where to run this command, see the topic on using the command-line tools.

Syntax

The command syntax is as follows:

```
restoreJobManager <location location name>
```

Parameters

The following parameter is available for the restoreJobManager command:

-location

Specifies the location of the backup database, including the database name. The parameter is required.

Usage scenario

The following example demonstrates correct syntax:

```
restoreJobManager -location /temp/jmbakup/OTiS
```

Example: Security and the command line tools

If you want to enable WebSphere Application Server security, you need to provide the command line tools with authentication information.

Without authentication information, the command line tools receive an `AccessDenied` exception when you attempt to use them with security enabled. There are multiple ways to provide authentication data:

- Most command line tools support a `-username` and `-password` option for providing basic authentication data. Specify the user ID and password for an administrative user. For example, you can use a member of the administrative console users with operator or administrator privileges, or the administrative user ID configured in the user registry. The following example demonstrates the **stopNode** command, which specifies command line parameters:

```
stopNode -username adminuser -password adminpw
```

- You can place the authentication data in a properties file that the command line tools read. The default file for this data is the `sas.client.props` file in the `properties` directory for the current profile.

Using MVS console commands

This topic discusses the MVS console commands that you can use to start, stop, and modify the product nodes and servers.

Before you begin

This topic assumes a basic familiarity with MVS console commands. See the Information Center for the z/OS product for information about how to use MVS operator commands.

About this task

You can issue commands from the MVS console yourself, or optionally use standard console automation products to automate product operations. Products such as Netview are presented copies of messages that are to be displayed on the MVS console. These automation products can also enter commands into the system using a virtual MVS console as a source.

You can use the MVS console commands to start, stop, and modify the product nodes and servers.

Procedure

- Start the application server, the deployment manager, or the node agent.
- Stop the application server, the deployment manager, or the node agent.
- Modify Application Server operations.
- Monitor your Application Server through display commands.

Results

Depending on which commands you issued, you might have stopped or started the Application Server, stopped or started the deployment manager, stopped or started the node agent, modified the Application Server operations, or monitored your Application Server.

START command

The **START** command reads the configuration file for the specified server process and starts the server, deployment manager, node agent, administrative agent, or job manager.

Depending on the options you specify, you can launch a new Java virtual machine (JVM) application programming interface (API) to run the server process or write the launch command data to a file. You can run this command from the MVS console of a WebSphere Application Server installation or a WebSphere Application Server, Network Deployment installation.

The following examples demonstrate correct syntax:

```
START appserver_proc_name,JOBNAME=server_short_name,  
    ENV=cell_short_name.node_short_name.server_short_name
```



```
START dmgr_proc_name,JOBNAME=server_short_name,  
    ENV=cell_short_name.node_short_name.server_short_name  
START nodeagent_proc_name,JOBNAME=server_short_name,  
    ENV=cell_short_name.node_short_name.server_short_name  
START administrative_agent_proc_name,JOBNAME=server_short_name,  
    ENV=cell_short_name.node_short_name.server_short_name  
START job_manager_proc_name,JOBNAME=server_short_name,  
    ENV=cell_short_name.node_short_name.server_short_name
```

You must enter this command on a single line. It is split here for display purposes.

STOP command

The **STOP** command reads the configuration file for the specified server process and stops the server, deployment manger, node agent, administrative agent, or job manager.

Depending on the options you specify, you can launch a new Java virtual machine (JVM) application programming interface (API) to run the server process or write the launch command data to a file. You can run this command from the MVS console of a WebSphere Application Server installation or a WebSphere Application Server, Network Deployment installation.

The following examples demonstrate correct syntax:

```
STOP appserver_proc_name,JOBNAME=server_short_name  
STOP dmgr_proc_name,JOBNAME=server_short_name  
STOP nodeagent_proc_name,JOBNAME=server_short_name  
STOP administrative_agent_proc_name,JOBNAME=server_short_name  
STOP job_manager_proc_name,JOBNAME=server_short_name
```

Modify command

You can issue the modify command from the MVS console to dynamically modify product operations.

You can use the modify command to display the status of various server components and activities, including the following components:

- Active controllers
- Trace settings
- Servants
- Sessions
- Java virtual machine (JVM) heap
- Java trace

Use the following format when entering the modify command.

```
f <server>, parameters
```

server is the short name for the server to which the modify command is directed. This argument is required.

Some parameters are split on multiple lines for publishing purposes.

Parameters

The following parameters are valid for the modify command:

DISABLEFAILOVER, '<JNDI Name>'

Used to manually disable resource failover support for the resource identified by the supplied JNDI name. The supplied JNDI name identifies the primary resource, and it must be surrounded by single quotes to preserve its case.

ENABLEFAILOVER, '<JNDI Name>'

Used to manually enable resource failover support for the resource identified by the supplied JNDI name. The supplied JNDI name identifies the primary resource, and it must be surrounded by single quotes to preserve its case.

FAILOVER, '<JNDI Name>'

Used to manually route requests that are destined for the primary resource to its configured alternate resource. The supplied JNDI name identifies the primary resource, and it must be surrounded by single quotes to preserve its case. To failback requests to the primary resource, issue the modify command with the parameter, **FAILBACK, '<JNDI name>'**.

FAILBACK, '<JNDI Name>'

Used to manually route requests to the configured primary resource that is identified by the supplied JNDI name. The supplied JNDI name identifies the primary resource, and it must be surrounded by single quotes to preserve its case. This command automatically enables automatic failover. If automatic failover is not wanted after failback, issue the modify command with the parameter, **DISABLEFAILOVER, '<JNDI Name>'**.

Attention: The JNDI name input is case sensitive.

CANCEL

Used to cancel the specified server.

You can specify the following options for the **CANCEL** parameter:

ARMRESTART

Specifies whether you are using the application response management (ARM) agent and want ARM to restart the server after it terminates. If you do not specify the ARMRESTART option on the **CANCEL** parameter, ARM does not restart the server.

HELP Specifies that you want to see the help information that describes the CANCEL syntax.

gotcha: You cannot use the **CANCEL** parameter to cancel a cluster from the MVS console. You must individually cancel each of the servers that make up the cluster.

```
DPM, [IOP=[nnn] | HTTP=[nnn] | HTTPS=[nnn] | MDB=[nnn] | SIP=[nnn] | SIPS=nnn] | [OLA=nnn] INTERVAL=[nnn]
| dump_action={SVC DUMP|JAVACORE|HEAP DUMP|TRACEBACK|JAVAT DUMP|NONE|RESET} | clear_all | reset_all]
```

The *nnn* placeholder specifies, in seconds, the length of the time interval for a specific protocol type. The value for the *nnn* placeholder must be an integer 0 - 255, or reset. A value of 0 illustrates that the protocol has no interval and that DPM is disabled for requests that use that protocol. A value of reset illustrates that the control of the DPM interval is returned to the `dpm_interval` classification attribute in the workload classification file.

The DPM INTERVAL modify command overrides the `dpm_interval` classification attribute. The DPM dump_action modify command overrides the `dpm_dump_action` classification attribute and the `server_region_dpm_dump_action` server-wide property.

Specifying RESET on the dump action deactivates the dump action for the modify command. The settings for the `dpm_dump_action` classification attribute and the `server_region_dpm_dump_action` server-wide property are again in effect.

A value specified on the **INTERVAL** parameter sets the DPM interval for all six protocols to the same value.

The `clear_all` value sets the time intervals for all related request protocols to 0 and the dump action to None.

The `reset_all` value resets all the DPM intervals and the **dump_action** parameter. Control of the DPM intervals and the dump action is returned to the WLM classification file.

HEAPDUMP

Used to generate a JVM heap dump for the controller, and for all the servants that are associated with that controller.

JAVACORE

Used to generate a JVM core dump for the controller, and for all the servants that are associated with that controller.

JAVATDUMP

Used to generate a JVM-initiated TDUMP for all the servant address spaces. You can also use the `JAVA_DUMP_TDUMP_PATTERN` native environment variable to specify the TDUMP data set name.

MSGROUTE, {ERRORLOG|HARDCOPY|CONSOLE|NONE}{[=xxx] [, RESET] [, CLEAR]}

Used to dynamically alter the output location of message IDs.

The `xxx` placeholder is a comma delimited list of product message IDs for the application server to route. Messages can be routed to more than one location.

The first time that you issue the command for a location and a set of messages, the message list for the location is the message list of the first command. If the environment variable for the location has messages defined, those messages are appended to the message list of the location. When you later issue the command for the same location and another set of messages, the application server appends the messages of the subsequent command to the existing list of messages.

Specifying `ERRORLOG` routes the messages to the error log instead of the default location.

Specifying `HARDCOPY` routes the messages to the `SYSLOG` system log, also known as `hardcopy`, instead of the default location.

Specifying `CONSOLE` routes the messages to the MVS operator console instead of the default location.

Specifying `NONE` suppresses the routing of the messages. A value of `NONE` takes precedence over routing of messages to the error log, the system log, or the MVS operator console.

Specifying `RESET` with `ERRORLOG`, `HARDCOPY`, `CONSOLE` or `NONE` means that the messages that are routed or suppressed come from the `was.env` file instead of the `MSGROUTE` modify command.

- Specifying `RESET` with `ERRORLOG` routes messages specified on the `ras_message_routing_errorlog` environment variable to the error log instead of the default location.
- Specifying `RESET` with `HARDCOPY` routes the messages specified on the `ras_message_routing_hardcopy` environment variable to the `SYSLOG` system log instead of the default location.
- Specifying `RESET` with `CONSOLE` routes the messages specified on the `ras_message_routing_console` environment variable to the MVS operator console instead of the default location.
- Specifying `RESET` with `NONE` suppresses the routing of messages specified on the `ras_message_routing_hardcopy` environment variable.

Specifying `CLEAR` with `ERRORLOG`, `HARDCOPY`, `CONSOLE` or `NONE` means that the messages are routed to a default location or are no longer suppressed.

- Specifying `CLEAR` with `ERRORLOG` routes messages to the default location instead of the error log.
- Specifying `CLEAR` with `HARDCOPY` routes the messages to the default location instead of to the `SYSLOG` system log.
- Specifying `CLEAR` with `CONSOLE` routes the messages to the default location instead of to the MVS operator console.
- Specifying `CLEAR` with `NONE` cancels the suppression of messages.

RECLASSIFY[,FILE=['/path/to/file.xml']]

Used to dynamically reload the workload classification document without having to restart the server. The reloaded workload classification settings take effect immediately.

Specifying RECLASSIFY performs reclassification using the workload classification file specified on the `wlm_classification_file` variable.

Specifying a file name with the reclassify parameter such as `RECLASSIFY,FILE='/path/to/newfile.xml'` performs reclassification using the classification file specified. The value for the `wlm_classification_file` variable changes to the file specified on the FILE option. Any subsequent reclassify modify commands that use the `wlm_classification_file` variable reload the new file specified on the variable, not the original file.

Specifying `RECLASSIFY,FILE=` disables classification. The `wlm_classification_file` variable is set to null. Any subsequent reclassify modify commands that use the `wlm_classification_file` variable do not affect the server since the `wlm_classification_file` variable has no value.

Attention: If the new workload classification file cannot be loaded, then the application server discards the reloaded classification settings. The application server continues to run with the classification settings in effect before the reclassify modify command being issued.

Tip: Whenever the workload classification file loads successfully, the application server issues to the system log, SYSLOG, a message like the following message. The message is also displayed on the operator console in response to the `DISPLAY,WORK,CLINFO` display command.

```
BBOJ0129I: The /tmp/wlm4.class.xml workload classification file
was loaded at 2009/07/14 19:33:35.297 (GMT)
```

STACKTRACE

Used to log the Java stack traceback for all Java threads in the server to the error log.

SMF,REQUEST,{ON|OFF|RESET}

Used to control the creation of SMF 120-9 records.

The `SMF,REQUEST` modify command overrides the `SMF_request_activity_enabled` classification attribute and the `server_SMF_request_activity_enabled` server-wide property.

Specifying `ON` turns on the creation of SMF 120-9 records.

Specifying `OFF` turns off the creation of SMF 120-9 records.

Specifying `RESET` deactivates the modify command setting for the creation of SMF 120-9 records. The settings for the `SMF_request_activity_enabled` classification attribute and the `server_SMF_request_activity_enabled` server-wide property at the time of the reset are in effect.

SMF,REQUEST,CPU,{ON|OFF|RESET}

Used in the collection of data for the CPU Usage Breakdown Section of the SMF 120-9 record whenever that record is created.

The `SMF,REQUEST,CPU` modify command overrides the `SMF_request_activity_CPU_detail` classification attribute and the `server_SMF_request_activity_CPU_detail` server-wide property.

Specifying `ON` turns on the collection of the data.

Specifying `OFF` turns off the collection of the data. When this value is specified, records that are created have zeros for this section in the location triplets at the top of the record.

Specifying `RESET` deactivates the modify command setting for the collection of the data. The settings for the `SMF_request_activity_CPU_detail` classification attribute and the `server_SMF_request_activity_CPU_detail` server-wide property at the time of the reset are in effect.

SMF,REQUEST,SECURITY,{ON|OFF|RESET}

Used in the collection of data for the Security Data Section of the SMF 120-9 record whenever that record is created.

The SMF,REQUEST,SECURITY modify command overrides the SMF_request_activity_security classification attribute and the server_SMF_request_activity_security server-wide property.

Specifying ON turns on the collection of the data.

Specifying OFF turns off the collection of the data. When this value is specified, records that are created have zeros for this section in the location triplets at the top of the record.

Specifying RESET deactivates the modify command setting for the collection of the data. The settings for the SMF_request_activity_security classification attribute and the server_SMF_request_activity_security server-wide property at the time of the reset are in effect.

SMF,REQUEST,TIMESTAMPS,{ON|OFF|RESET}

Used in the collection of data for the Formatted Timestamps Section of the SMF 120-9 record whenever that record is created.

The SMF,REQUEST,TIMESTAMPS modify command overrides the SMF_request_activity_timestamps classification attribute and the server_SMF_request_activity_timestamps server-wide property.

Specifying ON turns on the collection of the data.

Specifying OFF turns off the collection of the data. When this value is specified, records that are created have zeros for this section in the location triplets at the top of the record.

Specifying RESET deactivates the modify command setting for the collection of the data. The settings for the SMF_request_activity_timestamps classification attribute and the server_SMF_request_activity_timestamps server-wide property at the time of the reset are in effect.

TIMEOUT_DELAY=n

Used to dynamically alter the timeout delay option variable, control_region_timeout_delay, to *n* seconds. This modify command is used with the TIMEOUTDUMPACTION modify command to set a timeout delay that enables the specified dump action, such as JAVACORE, to complete when a timeout occurs. The dump actions specified on the TIMEOUTDUMPACTION modify command require a timeout delay that enables the servant to complete the dump action before the servant terminates.

After the specified dump is processed, you can reissue this command to change the value of the control_region_timeout_delay variable back to its previous setting without recycling the server.

TIMEOUTDUMPACTION=n

Used to indicate which of the following actions is performed whenever a timeout occurs for work that has been dispatched to a servant when the control_region_timeout_delay property is set to a non-zero value:

- If NONE, or none is specified, no dump is taken.
- If JAVACORE or javacore is specified, a Java core dump is taken.
- If SVCDDUMP or svcdump is specified, an SVC dump is taken.
- If JAVATDUMP or javatdump is specified, a JVM-initiated TDUMP is taken.

Note: The control_region_timeout_delay property must be configured to a length of time that is long enough for the servant to collect the JAVATDUMP before the controller terminates the servant. The TIMEOUT_DELAY modify command can be used to dynamically change the value of the control_region_timeout_delay property.

TIMEOUTDUMPACTIONSESSION=n

Used to indicate which of the following actions is performed whenever a timeout occurs for an HTTP, HTTPS, SIP, or SIPS request that has been dispatched to a servant, and the corresponding recovery property is set to SESSION:

- If NONE, or none is specified, no dump is taken.
- If JAVACORE or javacore is specified, a Java core dump is taken.
- If SVCDDUMP or svcdump is specified, an SVC dump is taken.
- If JAVATDUMP or javatdump is specified, a JVM-initiated TDUMP is taken

Following is a list of the corresponding recovery properties:

- protocol_http_timeout_output_recovery
- protocol_https_timeout_output_recovery
- protocol_sip_timeout_output_recovery
- protocol_sips_timeout_output_recovery

TRACEALL=n

Used to establish a general trace level for the server.

The following values are valid trace levels. Typically, you specify a value of 1.

- 0: no tracing is performed
- 1: tracing is performed when an exception occurs
- 2: basic tracing is performed
- 3: detailed tracing for all components is performed.

gotcha: Be careful when using a level of 3 because this level of tracing might yield more data than is practical for you to review.

TRACEBASIC=n

Used to specify the product components for which you want to turn on a basic level of tracing.

This command can override a different tracing level established by TRACEALL for those components.

gotcha: Do not change this variable unless directed to do so by IBM Support.

The following table includes the values that you can specify for this parameter. You can specify one or more of these values for either TRACEBASIC or TRACEDETAIL:

Value	Product component
0	RAS
1	Common Utilities
3	COMM
4	ORB
6	OTS
7	Shasta
9	z/OS Wrappers
A	Daemon
E	Security
F	Externalization
J (Should be used only under the direction of IBM Support)	JRas
L	Java EE

TRACEDETAIL=n

Used to specify the product components for which you want to turn on a detailed level of tracing.

This command activates the most detailed tracing for the specified product components and overrides different settings inTRACEALL. The selected components are identified by their component IDs, which are the same IDs as the ones listed for the TRACEBASIC parameter. Subcomponents, specified by numbers, receive detailed traces. Other parts of the product receive tracing as specified on the TRACEALL parameter.

gotcha: Do not change this variable unless directed to do so by IBM Support.

TRACESPECIFIC=xyyyzzz

Used to specify tracing overrides for specific product trace points.

Trace points are specified by eight-digit, hexadecimal numbers. To specify more than one trace point, use parentheses and separate the numbers with commas. You can also specify an environment variable name by enclosing the name in single quotation marks. The value of the environment variable is processed as if you had specified that value on the TRACESPECIFIC parameter.

gotcha: Do not use TRACESPECIFIC unless directed by IBM Support.

TRACE_EXCLUDE_SPECIFIC=xyyyzzz

Used to specify product trace points to exclude.

Trace points to exclude are specified by eight-digit, hexadecimal numbers. To specify more than one trace point, use parentheses and separate the numbers with commas. You can also specify an environment variable name by enclosing the name in single quotation marks. The value of the environment variable is processed as if you had specified that value on the TRACE_EXCLUDE_SPECIFIC parameter. You can use the TRACE_EXCLUDE_SPECIFIC parameter as a mask to turn off traces that are currently turned on. For example, use the TRACESPECIFIC parameter to turn on tracing for a whole part of the product, and then use the TRACE_EXCLUDE_SPECIFIC parameter to turn off one trace within that part of the product.

gotcha: Do not use TRACE_EXCLUDE_SPECIFIC parameter unless you are directed to do so by IBM Support.

TRACEINIT

Used to reset to the initial trace settings.

TRACENONE

Used to turn off all trace settings.

TRACETOSYSPRINT={YES|NO}

Used to select whether to send the trace to SYSPRINT.

Specifying YES sends the trace to SYSPRINT, and specifying NO stops the sending of the trace to SYSPRINT.

TRACETOTRCFILE={YES|NO}

Used to specify whether to direct the trace to the TRCFILE DD card.

Specifying YES sends the trace to the TRCFILE DD card, and specifying NO stops the sending of the trace to the TRCFILE DD card.

TRACEJAVA

Modify the Java trace string.

The Java trace specification is used to control Java tracing and conforms to the Java trace specification rules. *=all=enabled means to enable all types of tracing for all registered trace components.

TRACERECORD, {ON|OFF|RESET}

Used to control the writing of trace records. The parameter is used with the classification_only_trace attribute defined in the workload classification file.

Specifying ON causes trace records to be written based on the current trace specification. Any classification_only_trace attributes are ignored.

Specifying OFF prevents trace records from being written.

CAUTION: Trace records based on the current trace specification or produced as part of normal runtime processing are not written.

Specifying RESET causes trace records to be written for the current trace specification as follows:

- If no `classification_only_trace` attributes are defined, then trace records are written for all processing threads. This situation is functionally equivalent to setting `TRACERECORD` to `0N`.
- If one or more `classification_only_trace` attributes are defined, then trace records are only written for processing threads whose classification contains a `classification_only_trace` attribute of `1`.

The default value is `RESET`.

HELP

Used to display a list of all the keywords that you can use with the `modify` command.

You can also use the `HELP` parameter after the `CANCEL`, and `DISPLAY` parameters to display lists of all the keywords you can use with either of these parameters.

PAUSELISTENERS

Used to prevent work from being accepted into the server. Use this parameter to shut down the communication listeners and purge any pending work in the work registry.

WLM_MIN_MAX=min,max

Used to dynamically modify the minimum and maximum number of servants settings (`wlm_minimumSRCount` and `wlm_maximumSRCount`). WLM dynamically adjusts the current number of active servants to fit within the new constraints. The new settings do not necessarily take effect immediately; WLM works toward the new minimum and maximum targets as quickly as possible, considering factors such as resource availability and workload.

gotcha: The `WLM_MIN_MAX` command is ignored for a server that is configured with a single server dynamic application environment, that is the `wlm_dynapplenv_single_server` property is set to `1`.

RESUMELISTENERS

Used to restart the communication listeners after issuing a `modify` command that includes the `PAUSELISTENERS` parameter. This parameter allows new work to be accepted into the server.

DISPLAY

Used to display the name of the server, the system name where the server is running, and the current code level.

You can specify the following options for this parameter:

- `SERVERS` displays the name of the server at which the command is directed, the system name, and the code level for each active server in the sysplex that is in the same cell.
- `SERVANTS`: displays a list of the address space identifiers (ASIDs) of the servants that are attached to the server against which you issued the `display` command.
- `TRACE` displays trace information for a server controller. You can further modify this command with one of the following options:
 - `SRS` displays trace information for all servants, one at a time.
 - `ALL` displays trace information for the controller and all servants one at a time.
 - `JAVA` displays the Java trace string settings for a server controller. You can further modify this command with one of the following options:
 - `SRS` displays Java trace information for all servants, one at a time.
 - `ALL` displays Java trace information for the controller and all servants one at a time.
 - `HELP` displays a list of all the keywords that you can use with the `modify display trace Java` command.
 - `HELP` displays a list of all the keywords that you can use with the `modify display trace` command.
- `JVMHEAP` displays the JVM heap information for a server controller. You can further modify this command with one of the following options:
 - `SRS` displays the JVM heap information for all servants, one at a time.
 - `ALL` displays the JVM heap information for the controller and all servants, one at a time.

- HELP displays a list of all the keywords that you can use with the modify display Javaheap command.
- LISTENERS displays the connection instance name, associated IP address, and listening port number. The associated IP address can display an asterisk (*) as a wildcard.
- CONNECTIONS displays each connection instance name and a count of the number of connections. Each connection instance is on a separate line. You can further modify this command with one of the following options:
 - NAME='name' displays the number of associated connections for the specified connection instance 'name'. If the connection name is located but has zero connections, the command returns a count of zero. If the connection name is not found, the command returns an error message.
 - LIST displays the remote host information for all the connections of each of the connection instances. If a connection instance name has no connections, the command returns only the connection instance name.
 - LIST, NAME='name' displays the remote host information for all connections of a specified connection instance 'name'.
- HELP displays a list of all the keywords that you can use with the modify command.
- ERRLOG displays the last 10 entries in the error log. You can further modify this option with one of the following options:
 - SRS displays the last 10 entries in the error log for all servants, one at a time.
 - HELP displays a list of all the keywords you can use with the modify display error log command.
- MODE displays the run bit mode.
- MSGROUTE displays the list of messages for all routing locations and for messages to suppress. Routing locations can be the error log, the system log, or the MVS operator console.
 - ERRORLOG displays a list of messages to route to the error log.
 - HARDCOPY displays a list of messages to route to the SYSLOG system log, also known as hardcopy.
 - CONSOLE displays a list of messages to route to the MVS operator console.
 - NONE displays a list of messages to suppress.
- SMF displays the status of SMF 120-9 record creation for a server.
- WORK displays the work element counts for the server. You can further modify this option with one of the following options:
 - EJB displays Enterprise JavaBeans (EJB) request count information. You can modify this option with one of the following options:
 - SRS displays EJB request count information for all servants, one at a time.
 - HELP displays a list of all the keywords you can use with the modify display work EJB command.
 - SERVLET displays servlet request count information. You can modify this option with one of the following options:
 - SRS displays servlet request count information for all servants, one at a time.
 - HELP displays a list of all the keywords you can use with the modify display work servlet command.
 - MDB displays message-driven bean (MDB) request count information. You can further modify this command with one of the following options:
 - SRS displays MDB request count information for all servants, one at a time.
 - HELP displays a list of all the keywords you can use with the modify display work MDB command.
 - SIP displays Session Initiation Protocol (SIP) request count information. You can modify this option with one of the following options:
 - SRS displays SIP request count information for all servants, one at a time.

- HELP displays a list of all the keywords you can use with the modify display work SIP command.
- SUMMARY displays summary request count information.
 - SRS displays summary request count information for all servants, one at a time.
 - HELP displays a list of all the keywords you can use with the modify display work summary command.
- ALL displays all request count information.
 - SRS displays all request count information for all servants, one at a time.
 - STATS displays all detailed statistics.
 - HELP displays a list of all the keywords you can use with the modify display work all command.
- CLINFO displays work classification information.
- HELP displays a list of all the keywords you can use with the modify display work command.
- OLA displays inbound Optimized Local Adapter (OLA) requests driven by services BBOA1INV and BBOA1SRQ. You can further modify this command with one of the following options:
 - SRS displays OLA request count information for all servants, one at a time.
 - HELP displays a list of all the keywords you can use with the modify display work all command.
- Fast Response Cache Accelerator (FRCA). Display FRCA cache statistics as follows:
 - f <servername>,display,frca
 - f <servername>,display,frca,content
 - f <servername>,display,frca,stats
- ADAPTER
 - STATUS
 - Displays the status of the optimized local adapters component.
 - REGISTRATIONS
 - Displays a list of optimized local adapters clients connected to this server.
 - DAEMONRGES
 - Displays a list of optimized local adapters clients connected to any server in this daemon group.
 - SEARCHRGES
 - Displays a list of optimized local adapters clients connected to any server in this daemon group, matching search criteria.

Example: Disabling fail over

Example 1: The following command disables fail over for the primary resource identified by jdbc/MySampleDataSource on the server bbos001:

```
f bbos001,DISABLEFAILOVER,'jdbc/MySampleDataSource'
```

Example: Canceling application clusters and servers with the modify command

You cannot cancel a cluster from the MVS console. Instead, you must cancel each of the servers that make up the cluster.

Example 1: The following command cancels the bbo6acr server:

```
f bbo6acr,cancel
```

Example 2: The following command cancels the bbo6acr server and instructs ARM to restart the bbo6acr server after the terminate process completes:

```
f bbo6acr,cancel,armrestart
```

Example: Establishing a general level of trace using scripting

To establish a general trace level for the server, use the following command:

```
f server,traceall=n
```

Valid trace levels are 0 (none), 1 (exception), 2 (basic), and 3 (detailed tracing). Under normal conditions and in production, use 1 (exception).

Example: The following command turns on exception level tracing for the *bbo6acr* server:

```
f bbo6acr,traceall=1
```

Here is a sample display:

```
f bbo6acr,traceall=1  
BB000211I MODIFY COMMAND TRACEALL=1 COMPLETED SUCCESSFULLY
```

Example: Setting basic and detailed trace levels for the WebSphere z/OS components

Use the following command:

```
f server,tracebasic=(n,...)
```

Example 1: The following command turns on a basic level of tracing for the Daemon component on server *bbo6acr*:

```
f bbo6acr,tracebasic=a
```

Here is a sample display:

```
f bbo6acr,TRACEBASIC=A  
BB000211I MODIFY COMMAND TRACEBASIC=A COMPLETED SUCCESSFULLY
```

To specify the WebSphere for z/OS components for which you want to switch on a detailed level of tracing use the following command:

```
f server,tracedetail=(n,...)
```

Example 2: The following command turns on a detailed level of tracing for security on server *bbo6acr*:

```
f bbo6acr,tracedetail=e
```

Here is a sample display:

```
f bbo6acr,TRACEDETAIL=E  
BB000211I MODIFY COMMAND TRACEDETAIL=E COMPLETED SUCCESSFULLY
```

Example: Setting specific trace points using scripting

To set specific trace points, use the following command:

```
f server,tracespecific=n | (n,...)
```

Example 1: The following command turns on the specific trace point 04006001:

```
f bbo6acr,tracespecific=04006001
```

Example 2: The following command turns on the specific trace points 04006001 and 04006027:

```
f bbo6acr,tracespecific=(04006001,04006027)
```

Example 3: The following command turns on the specific trace points set in the environment variable *tracepoints*:

```
f bbo6acr,tracespecific='tracepoints'
```

Example: Excluding specific trace points using scripting

To exclude specific trace points, use the following command:

```
f server,trace_exclude_specific=n | (n,...)
```

Example: The `tracespecific` command turns on tracing for a whole part and then the `trace_exclude_specific` turns off tracing for the point 04006031 within that part:

```
f bbo6acr,tracespecific=04006000
```

```
f bbo6acr,trace_exclude_specific=04006031
```

Example: Resetting to the initial trace settings using scripting

To reset to the initial trace settings use the following command:

```
f server,traceinit
```

Example: Command to turn off tracing using a command line

To turn off tracing, use the following command:

```
f server,tracenone
```

Example: Sending the trace to SYSPRINT using scripting

To send the trace to SYSPRINT use the following command:

```
f server,tracetosysprint=yes
```

To stop sending the trace to SYSPRINT, use the following command:

```
f server,tracetosysprint=no
```

Example: Canceling application clusters and servers with the modify command

This example demonstrates how to cancel application clusters and servers with the `modify` command.

Before you begin: You cannot cancel a cluster from the MVS console. Instead, you must cancel each of the servers that make up the cluster.

Example 1: The following command cancels the *bbo6acr* server:

```
f bbo6acr,cancel
```

Example 2: The following command cancels the *bbo6acr* server and instructs ARM to restart it after it terminates:

```
f bbo6acr,cancel,armrestart
```

Example: Getting help for the modify command

You can get syntax help for the various levels of the `modify` command.

Example 1: Enter the following command to display a list of all the keywords that you can use with the `modify` command:

```
f bbo6acr,help
```

The resulting output is like the following example:

```

F bbo6acr,HELP
BB000178I THE COMMAND MODIFY MAY BE FOLLOWED BY ONE OF THE FOLLOWING KEYWORDS:
BB000179I CANCEL - CANCEL THIS CONTROL REGION
BB000179I TRACEALL - SET OVERALL TRACE LEVEL
BB000179I TRACEBASIC - SET BASIC TRACE COMPONENTS
BB000179I TRACEDetail - SET DETAILED TRACE COMPONENTS
BB000179I TRACESPECIFIC - SET SPECIFIC TRACE POINTS
BB000179I TRACEINIT - RESET TO INITIAL TRACE SETTINGS
BB000179I TRACENONE - TURN OFF ALL TRACING
BB000179I TRACETOSYSPRINT - SEND TRACE OUTPUT TO SYSPRINT (YES/NO)
BB000179I TRACETOTRCFILE - SEND TRACE OUTPUT TO TRCFILE DD CARD (YES/NO)
BB000179I DISPLAY - DISPLAY STATUS
BB000179I TRACE_EXCLUDE_SPECIFIC - EXCLUDE SPECIFIC TRACE POINTS
BB000179I TRACEJAVA - SET JAVA TRACE OPTIONS
BB000179I TRACETOTRCFILE - SEND TRACE OUTPUT TO TRCFILE (YES/NO)
BB000179I MDBSTATS - MDB DETAILED STATISTICS
BB000179I PAUSELISTENERS - PAUSE THE COMMUNICATION LISTENERS
BB000179I RESUMELISTENERS - RESUME THE COMMUNICATION LISTENERS
BB000179I STACKTRACE - LOG JAVA THREAD STACK TRACEBACKS
BB000179I TIMEOUTDUMPACTION - SET TIMEOUT DUMP ACTION
BB000179I TIMEOUTDUMPACTIONSESSION - SET TIMEOUT DUMP ACTION SESSION
BB000179I WLM_MIN_MAX - RESET WLM MIN/MAX SERVANT SETTINGS

```

Example 2: Enter the following command to display a list of all the keywords you can use with the modify display command:

```
f bbo6acr,display,help
```

The resulting output is similar to the following example:

```

F BBOS001,DISPLAY,HELP
BB000178I THE COMMAND DISPLAY, MAY BE FOLLOWED BY ONE OF THE FOLLOWING
124
KEYWORDS:
BB000179I SERVERS - DISPLAY ACTIVE CONTROL PROCESSES
BB000179I SERVANTS - DISPLAY SERVANT PROCESSES OWNED BY THIS CONTROL 126
PROCESS
BB000179I LISTENERS - DISPLAY LISTENERS
BB000179I CONNECTIONS - DISPLAY CONNECTION INFORMATION
BB000179I TRACE - DISPLAY INFORMATION ABOUT TRACE SETTINGS
BB000179I JVMHEAP - DISPLAY JVM HEAP STATISTICS
BB000179I WORK - DISPLAY WORK ELEMENTS
BB000179I ERRLOG - DISPLAY THE LAST 10 ENTRIES IN THE ERROR LOG
BB000179I MODE - DISPLAY THE EXECUTION BITMODE
BB000188I END OF OUTPUT FOR COMMAND DISPLAY,HELP

```

Example 3: Enter the following command to display a list of all the keywords that you can use with the modify timeoutdumpaction or modify timeoutdumpactionsession command:

```
f bbos001,timeoutdumpactionsession=help
```

The resulting output is like the following example:

```

BB000178I THE COMMAND MODIFY TIMEOUTDUMPACTIONSESSION= MAY BE FOLLOWED BY
ONE OF THE FOLLOWING KEYWORDS:
BB000179I SVCDUMP - SVC DUMP
BB000179I JAVACORE - JAVA CORE DUMP
BB000179I NONE - NO DUMP

```

Example 4: Enter the following command to display a list of all the keywords that you can use with the modify display dpm command:

```
f bbos001,dpm,help
```

The resulting output is like the following example:

```
SY1 BB000178I THE COMMAND MODIFY DPM, MAY BE FOLLOWED BY ONE OF THE
FOLLOWING KEYWORDS:
BB000179I IIOP - IIOP DISPATCH PROGRESS MONITOR INTERVAL
BB000179I HTTP - HTTP DISPATCH PROGRESS MONITOR INTERVAL
BB000179I HTTPS - HTTPS DISPATCH PROGRESS MONITOR INTERVAL
BB000179I MDB - MDB DISPATCH PROGRESS MONITOR INTERVAL
BB000179I SIP - SIP DISPATCH PROGRESS MONITOR INTERVAL
BB000179I SIPS - SIPS DISPATCH PROGRESS MONITOR INTERVAL
BB000179I INTERVAL - ALL DISPATCH PROGRESS MONITOR INTERVALS
BB000179I DUMP_ACTION - DISPATCH PROGRESS MONITOR DUMP ACTION
BB000179I CLEAR_ALL - CLEAR ALL DISPATCH PROGRESS MONITOR SETTINGS
BB000179I RESET_ALL - RESET ALL DISPATCH PROGRESS MONITOR SETTINGS
BB000211I MODIFY COMMAND DPM, COMPLETED SUCCESSFULLY
```

Example: Modifying the Java trace string

This example shows how to use a console command to modify a Java trace string.

To modify the Java trace string, use the following command:

```
f server, tracejava='trace specification'
```

Example: The following command enables all types of tracing for all registered trace components:

```
f bbo6acr, tracejava='*=all'
```

Note: The quotation marks are a required part of the syntax.

Display command with examples

The display command is useful for obtaining data on a variety of processes.

Use the following display command examples to monitor how the product is functioning.

- “Example: Displaying active replies” on page 180
- “Example: Displaying active address spaces” on page 180
- “Example: Displaying the status of address spaces registered for automatic restart management (ARM)” on page 180
- “Example: Displaying units of work (transactions) for the Information Management System” on page 180
- “Example: Displaying JVM heap information” on page 181
- “Example: Displaying status of a server” on page 182
- “Example: Displaying the status of SMF 120-9 record creation” on page 182
- “Example: Displaying the status of dispatch thread” on page 183
- “Example: Displaying status of clusters” on page 185

How to display servants using scripting.

To display servants, use the following command:

```
f server, display, servants
```

Example: The following command will display servants for the bbo6acr server:

```
f bbo6acr, display, servants
```

Here is a sample display:

```
F bbo6acr, DISPLAY, SERVANTS
BB000185I SERVER BB05SR4/BB0ASR4A HAS 1 SERVANT PROCESS (ASID: 0038x)
BB000188I END OF OUTPUT FOR COMMAND DISPLAY, SERVANTS
```

How to display trace settings and Java string trace settings with scripting.

To display the trace settings for a server instance, use the following command:

```
f server,display,trace
```

Example 1: The following command will display trace settings for the bbo6acr server:

```
f bbo6acr,display,trace
```

Here is a sample display:

```
F bbo6acr,DISPLAY,TRACE
BB000224I TRACE INFORMATION FOR SERVER BB05SR4/BB0ASR4A/STC00047
BB000197I LOCATION = SYSPRINT BUFFER
BB000197I AGGREGATE TRACE LEVEL = 1
BB000197I EXCEPTION TRACING = RAS(0), Common Utilities(1), COMM(3),
ORB(4), OTS(6), Shasta(7), z/OS Wrappers(9), Daemon(A), Security(E),
Externalization(F), JRAS(J), J2EE(L)
BB000197I BASIC TRACING =
BB000197I DETAILED TRACING =
BB000197I TRACE SPECIFIC = NONE SPECIFIED
BB000197I TRACE EXCLUDE SPECIFIC = NONE SPECIFIED
BB000225I TRACE INFORMATION FOR SERVER BB05SR4/BB0ASR4A/STC00047 COMPLETE
BB000188I END OF OUTPUT FOR COMMAND DISPLAY,TRACE
```

To display the Java trace string settings for a server, use the following command:

```
f server,display,trace,java
```

Example 2: The following command will display Java trace settings for the bbo6acr server:

```
f bbo6acr,display,trace,java
```

Here is a sample display:

```
F bbo6acr,DISPLAY,TRACE,JAVA
BB000196I TRACE INFORMATION FOR SERVER BB05SR4/BB0ASR4A
BBOJ0050I CTL(STC00047):*=all=disabled
BB000188I END OF OUTPUT FOR COMMAND DISPLAY,TRACE,JAVA
```

To display the Java trace information for all servants, one at a time.

```
f server,display,trace,java.srs
```

Example 3: The following command will display Java trace settings for the bbo6acr server:

```
f bbo6acr,display,trace,java,srs
```

Here is a sample display:

```
F bbo6acr,DISPLAY,TRACE,JAVA,SRS
BB000196I TRACE INFORMATION FOR SERVER BB05SR4/BB0ASR4A
BBOJ0050I SR(STC00048):*=all=disabled
BB000188I END OF OUTPUT FOR COMMAND DISPLAY,TRACE,JAVA,SRS
```

To display the Java trace information for the controller and all servants, one at a time.

```
f server,display,trace,java.all
```

Example 4: The following command will display Java trace settings for the controller and all servants:

```
f bbo6acr,display,trace,java,all
```

Here is a sample display:

```
F bbo6acr,DISPLAY,TRACE,JAVA,ALL
BB000196I TRACE INFORMATION FOR SERVER BB05SR4/BB0ASR4A
BB0J0050I CTL(STC00047):*=all=disabled
BB0J0050I SR(STC00048):*=all=disabled
BB000188I END OF OUTPUT FOR COMMAND DISPLAY,TRACE,JAVA,ALL
```

Example: Displaying active replies

Displaying active replies from the MVS console allows you to observe system activity and determine if the system requires an operator response.

Issue the following command to display (list) all active replies:

```
d r,r
```

Example: Displaying active address spaces

Use the display command to display active address spaces.

For example, use the display command to determine if the location service daemon is up:

Table 12. display command descriptions. The following table describes the display command.

d a,l	Displays a list of all address spaces.
d a,a	Displays a list of all active address spaces.
d a,address-space-name	Displays only the address space in which you are interested. This command is recommended over the first two because it does not yield such a lengthy list on a production system. Of course, you need to know the name of the address space for which you are looking. Example: d a,bboasr1
d a,bbo*	Displays a list of all active address spaces that start with BB0.

Example: Displaying the status of address spaces registered for automatic restart management (ARM)

This example shows how to display the status of addresses registered with automatic restart management.

Perform the following steps to use automatic restart management (ARM) to display the status of ARM registered address spaces (including the address spaces of server instances) in the WebSphere Application Server for z/OS environment:

1. Initialize all servers.
2. Issue one or both of the following commands:
 - To display all registered address spaces (including the address spaces of server instances), issue the command:
d xcf,armstatus,detail
 - To display the status of a particular server instance, use the display command and identify the job name. For example, to display the status of the Daemon server instance (job BBODMN), issue the following command:
d xcf,armstatus,jobname=bbodmn,detail

Example: Displaying units of work (transactions) for the Information Management System

This example shows how to use scripting to display the units of work for the Information Management System.

You can display units of work (transactions) for Information Management System (IMS™):

1. To display the status of a specific transaction, issue this command. The message ID is the replying IMS ready message number.

```
message_id/dis tran trans-name
```

2. To display the status of a specific program, issue the command:

```
/dis prog program-name
```

3. To display the number of Message Processing Regions (MPRs) that are currently active, issue the command:

```
/display active region
```

For more information about IMS commands, see the *IMS/ESA® Summary of Operator Commands* manual on the IBM Publications Center website.

Example: Displaying JVM heap information

This example demonstrates how to display JVM heap information with scripting.

To display the JVM heap information for a server controller, use the following command:

```
f server,display,jvmheap
```

Example 1: The following command will display JVM heap information for the bbo6acr server:

```
f bbo6acr,display,jvmheap
```

Here is a sample display:

```
F bbo6acr,DISPLAY,JVMHEAP
BB000201I JVM HEAP INFORMATION FOR SERVER BB05SR4/BBOASR4A/STC00047
BB000202I (STC00047) HEAP(MIDDLEWARE), COUNT(00000000), FREE STORAGE(
396FA70), TOTAL STORAGE( 7FFFA00)
BB000204I JVM HEAP INFORMATION FOR SERVER BB05SR4/BBOASR4A/STC00047 COMPLETE
BB000188I END OF OUTPUT FOR COMMAND DISPLAY,JVMHEAP
```

To display the JVM heap information for all servants one at a time, use the following command:

```
f server,display,jvmheap,srs
```

Example 2: The following command will display JVM heap information for all servants in the bbo6acr server:

```
f bbo6acr,display,jvmheap,srs
```

Here is a sample display:

```
F bbo6acr,DISPLAY,JVMHEAP,SRS
+BB000201I JVM HEAP INFORMATION FOR SERVER BB05SR4/BBOASR4A/STC00048
+BB000202I (STC00048) HEAP(MIDDLEWARE), COUNT(00000001), FREE STORAGE(
 25F4030), TOTAL STORAGE( 7FFFA00)
+BB000204I JVM HEAP INFORMATION FOR SERVER BB05SR4/BBOASR4A/STC00048 COMPLETE
BB000188I END OF OUTPUT FOR COMMAND DISPLAY,JVMHEAP,SRS
```

To display the JVM heap information for the controller and all servants of a server one at a time, use the following command:

```
f server,display,jvmheap,all
```

Example 3: The following command will display JVM heap information for the controller and all servants in the bbo6acr server:

```
f bbo6acr,display,jvmheap,all
```

Here is a sample display:

```

F bbo6acr,DISPLAY,JVMHEAP,ALL
BB000201I JVM HEAP INFORMATION FOR SERVER BB05SR4/BBOASR4A/STC00047
BB000202I (STC00047) HEAP(MIDDLEWARE), COUNT(00000000), FREE STORAGE(
396FA70), TOTAL STORAGE( 7FFFA00)
BB000204I JVM HEAP INFORMATION FOR SERVER BB05SR4/BBOASR4A/STC00047 COMPLETE
+BB000201I JVM HEAP INFORMATION FOR SERVER BB05SR4/BBOASR4A/STC00048
+BB000202I (STC00048) HEAP(MIDDLEWARE), COUNT(00000001), FREE STORAGE(
 25F4030), TOTAL STORAGE( 7FFFA00)
+BB000204I JVM HEAP INFORMATION FOR SERVER BB05SR4/BBOASR4A/STC00048 COMPLETE
BB000188I END OF OUTPUT FOR COMMAND DISPLAY,JVMHEAP,ALL

```

Example: Displaying status of a server

This example shows how to use scripting to display the status of a server.

To display the status of a server, use the following command:

```
f server,display
```

Example: The following command will display status for the bbo6acr server:

```
f bbo6acr,display
```

Here is a sample display:

```

F bbo6acr,DISPLAY
BB000173I SERVER BB05SR4/BBOASR4A ACTIVE ON SY1 AT LEVEL wd5xo03.
BB000188I END OF OUTPUT FOR COMMAND DISPLAY

```

Example: Displaying the status of SMF 120-9 record creation

The SMF display command displays the status of SMF 120-9 record creation and whether CPU usage, security information, or, timestamp information for the record is collected.

The display,smf command also displays the following information:

- When the last SMF 120-9 record was successfully or unsuccessfully created
- How many SMF records were successfully and unsuccessfully created since the server started
- The last non-zero return code from an attempt to write an SMF 120-9 record

To display the status of SMF 120-9 record creation for a server, use the following command:

```
f server,display,smf
```

Example: The following command displays the status of SMF 120-9 record creation for the bbo6acr server:

```
f bbo6acr,display,smf
```

Here is a sample display:

```

BB000344If
  SMF 120-9: ON, CPU USAGE: OFF, TIMESTAMPS: OFF, SECURITY INFO: OFF
BB000345I SMF 120-9: TIME OF LAST WRITE: 2008/02/20 13:38:44.003116, SUCCESSFUL
  WRITES: 2, FAILED WRITES: 0
BB000346I SMF 120-9: LAST FAILED WRITE TIME: NEVER, RC: 0
BB000188I END OF OUTPUT FOR COMMAND DISPLAY,SMF

```

Each of the SubType 9, CPU Usage, Timestamps, or Security Info sections of the message can have a value of ON, OFF, FORCED_ON, or FORCED_OFF. The ON or OFF values for a section indicate the setting of the corresponding server-wide custom property. If any values are FORCED_ON or FORCED_OFF, then an SMF modify command was issued. The SMF modify command overrides the SMF server-wide custom property and the SMF request-level RAS attribute.

Example: Displaying the status of dispatch thread

This example shows how to display the status of dispatch threads. You can display the status of each servant thread, on which there is currently a request dispatched, for all of the servants in a controller, or you can display this thread status for a particular servant within the controller. You can also display information about the dispatch thread for a particular request, the dispatch threads that have been handling the same request for a specific amount of time, or the dispatch threads for requests that have timed out.

To display the status of dispatch threads use the following command:

```
f server,display,threads,<parameters>
```

server is the short name for the server to which the this command is directed. This argument is required.

If you do not specify any parameters, all is used.

ALL

Used to display information about all of the requests that are currently dispatched in each servant that is associated with this controller.

You can specify the following options on the DISPLAY,THREADS,ALL command. If no option is specified, summary is used.

DETAILS

Specifies that you want to see the full details about each request.

SUMMARY

Specifies that, for each request, you only want to see the request ID, the ASID of the servant where the request was dispatched, the time the request was dispatched, and whether the request timed out.

Here is a sample display of the messages that are issued if you enter this command:

Example: The following command displays the summary status of all of the dispatch threads for the BBOS001 server:

```
f BBOS001,display,threads,all
```

The following messages display in response to this command. In this output, the requests that have a Y in the JW column were in a Java wait state when the command was issued. Similarly, the requests that have a Y in the TO column were considered unresponsive.

As part of its request processing, the product uses several techniques in attempts to make an unresponsive thread responsive, and enable a request to complete. However, these attempts are not always successful. After the product exhausts the maximum number of attempts that it can make to get a request to complete, it notifies the controller that the request should be terminated. In the preceding example, the requests with a Y in the RE column are the requests for which the product has exhausted attempts to get the request to finish.

```
BBOJ0111I: REQUEST ASID JW TO RE DISPATCH TIME
BBOJ0112I: fffffb41d 0X0176 N N N 2008/03/24 20:03:28.212400
BBOJ0112I: fffffb420 0X0176 Y Y Y 2008/03/24 19:59:16.910228
BBOJ0112I: fffffb421 0X0176 Y Y Y 2008/03/24 19:59:14.982621
BBOJ0112I: fffffb422 0X0176 Y Y Y 2008/03/24 19:59:12.805119
BBOJ0112I: fffffb423 0X0176 Y Y Y 2008/03/24 19:59:08.936107
BBO00188I END OF OUTPUT FOR COMMAND DISPLAY,THREADS,ALL
```

Example: The following command displays the detailed status of all of the dispatch threads for the BBOS001 server. For this command, the ALL parameter is implied.

```
f BBOS001,display,threads,details
```

The following messages display in response to this command:

```

BBOJ0106I: REQUEST fffff2a9 ASID 0X0064 TCB 0X008D0438
BBOJ0119I: CONTROLLER RECEIVED REQUEST AT 2008/06/05 14:24:51.508941
BBOJ0120I: CONTROLLER QUEUED REQUEST TO WLM AT 2008/06/05
14:24:51.509147
BBOJ0107I: SERVANT DISPATCHED REQUEST AT 2008/06/05 14:24:51.509242
BBOJ0108I: JVM THD IS NOT HUNG: ITI INACTIVE
BBOJ0110I: DETAILS FOR GIOP Outbound: Target Operation =sayHelloTwo
BBOJ0110I: DETAILS FOR JVM INTERRUPTIBLE THREAD: Monitor ACTIVE
BBOJ0106I: REQUEST fffff2a3 ASID 0X0064 TCB 0X008CB2D8
BBOJ0119I: CONTROLLER RECEIVED REQUEST AT 2008/06/05 14:24:54.981645
BBOJ0120I: CONTROLLER QUEUED REQUEST TO WLM AT 2008/06/05
14:24:54.982324
BBOJ0107I: SERVANT DISPATCHED REQUEST AT 2008/06/05 14:24:55.522788
BBOJ0108I: JVM THD IS NOT HUNG: ITI INACTIVE
BBOJ0110I: DETAILS FOR JVM INTERRUPTIBLE THREAD: Monitor ACTIVE
BBO00188I END OF OUTPUT FOR COMMAND DISPLAY,THREADS,DETAILS

```

TIMEDOUT

Used to display information about the requests that are currently considered timed out.

You can specify the following options for the DISPLAY,THREADS,TIMEDOUT command. If no option is specified, summary is used.

DETAILS

Specifies that you want the full details about the requests that timed out.

SUMMARY

Specifies that, for each request that timed out, you only want to see the request ID, the ASID of the servant where the request was dispatched, and the time that the request was dispatched

Example: The following command displays a summary of the requests, that are running on the BBOS001 server, that are currently considered timed out.

```
f BBOS001,display,threads,timedout,summary
```

REQUEST=<value>

Used to display information about the request that has the specified ID.

You can specify the following options on the DISPLAY,THREADS,REQUEST command. If no option is specified, details is used.

DETAILS

Specifies that you want to see the full details about this request

SUMMARY

Specifies that, for this request, you only want to see the ASID of the servant where the specified request was dispatched, the time the request was dispatched, and whether the request timed out.

Example: The following command displays a summary of the request FFFFF2A9. For this command, the DETAILS option is implied.

```

f BBOS001,display,threads,request=FFFFF2A9
BBOJ0106I: REQUEST fffff2a9 ASID 0X0064 TCB 0X008D0438
BBOJ0119I: CONTROLLER RECEIVED REQUEST AT 2008/03/27 16:51:36.981645
BBOJ0120I: CONTROLLER QUEUED REQUEST TO WLM AT 2008/03/27
16:51:36.982324
BBOJ0107I: SERVANT DISPATCHED REQUEST AT 2008/03/27 16:51:36.987541
BBOJ0108I: JVM THD IS NOT HUNG: ITI INACTIVE
BBOJ0110I: DETAILS FOR GIOP Outbound: Target Operation = sayHelloTwo
BBOJ0110I: DETAILS FOR JVM INTERRUPTIBLE THREAD: Monitor ACTIVE
BBO00188I END OF OUTPUT FOR COMMAND DISPLAY,THREADS,REQUEST=FFFFF2A9

```

ASID=<value>

Used to display information about the request, with the specified ASID, that is currently dispatched in the servant. The ASID must be specified as a hexadecimal value.

You can specify the following options on the DISPLAY,THREADS,ASID command. If no option is specified, summary is used.

DETAILS

Specifies that you want to see the full details about these requests.

SUMMARY

Specifies that you only want to see the request IDs, the time that each request was dispatched, and whether the requests timed out.

Example: The following command displays a summary of the dispatch threads for the servant that has the ASID 0038x.

```
f BB0S001,display,threads,ASID=0038x
```

AGE=<value>

Used to display information about the requests, for which the amount of time that has elapsed since the request was placed in dispatch is equal to or greater than the specified amount of time. Specify the length of time in seconds.

You can specify the following options on the DISPLAY,THREADS,AGE command. If no option is specified, summary is used.

DETAILS

Specifies that you want to see the full details about these requests.

SUMMARY

Specifies that, for each request, you only want to see the request ID, ASID where the request was dispatched, the time the request was dispatched, and whether the request timed out.

Example: The following command displays a summary of the dispatch threads for the requests for which the amount of time, that has elapsed since the request was placed in dispatch, is equal to or greater than 30 seconds

```
f BB0S001,display,threads,age=30
```

Example: Displaying status of clusters

This example shows how to use scripting to display the status of clusters.

To display the status of clusters in the active controller, use the following command:

```
f server,display,servers
```

Example: The following command will display status for the bbo6acr server:

```
f bbo6acr,display,servers
```

Here is a sample display:

```
F bbo6acr,DISPLAY,SERVERS
BB000182I SERVER          ASID  SYSTEM  LEVEL
BB000183I CBDAEMON/DAEMON01  31x  SY1     wd5xo03
BB000183I BB05SR4 /BBOASR4A  1F6x SY1     wd5xo03
BB000188I END OF OUTPUT FOR COMMAND DISPLAY,SERVERS
```

Example: Displaying units of work (transactions) for the Information Management System

This example shows how to use scripting to display the units of work for the Information Management System.

You can display units of work (transactions) for Information Management System (IMS):

1. To display the status of a specific transaction, issue this command. The message ID is the replying IMS ready message number.

message_id/dis tran trans-name

2. To display the status of a specific program, issue the command:

/dis prog program-name

3. To display the number of Message Processing Regions (MPRs) that are currently active, issue the command:

/display active region

For more information about IMS commands, see the *IMS/ESA Summary of Operator Commands* manual on the IBM Publications Center website.

Example: Displaying the setting for trace records

Use the following example to display the setting for trace records. The TRACERECORD parameter on the modify command is used to control the writing of trace records.

To display the setting of the TRACERECORD parameter for a server, use the following command:

```
F server,DISPLAY,TRACERECORD
```

Example: The following command displays the setting of the TRACERECORD parameter for the bbo6acr server:

```
F bbo6acr1,DISPLAY,TRACERECORD
```

See the following sample display:

```
F BB06ACR,DISPLAY,TRACERECORD
BB000365I THE TRACERECORD PARAMETER IS ON.
BB000188I END OF OUTPUT FOR COMMAND DISPLAY,TRACERECORD
```

Chapter 4. Using Ant to automate tasks

To support using Apache Ant with Java Platform, Enterprise Edition (Java EE) applications running on the application server, the product provides a copy of the Ant tool and a set of Ant tasks that extend the capabilities of Ant to include product-specific functions. Ant has become a very popular tool among Java programmers.

About this task

Apache Ant is a Java-based build tool. In theory, it is similar to Make, but Ant is different. Instead of a model in which it is extended with shell-based commands, Ant is extended using Java classes. Instead of writing shell commands, XML-based configuration files are used. These files reference a target tree in which various tasks are run. Each task is run by an object that implements a particular Task interface.

gotcha:

- Calling the WebSphere_Ant (WsAnt) scripts outside of the ws_ant launcher is not supported or recommended.
- When you invoke the Ant tool, do not pass empty strings in place of command arguments in ant script. The script will not work in the wsadmin environment. The `<arg value>` cannot be an empty string, such as shown in the following example:

```
<?xml version="1.0" encoding="UTF-8"?>
<project name="proj" default="main">
  <taskdef name="wsadmin" classname="com.ibm.websphere.ant.tasks.WsAdmin"/>
  <target name="main">
    <wsadmin,conntype="NONE" lang="jython" failonerror="true" script="&(basedir)/script.ph">
      <arg value="blah" />
      <arg value="" />
    </wsadmin>
    <!-- manaeapp action="blah" variation-number="0" -->
  </target>
</project>
```

By combining the following tasks with those provided by Ant, you can create build scripts that compile, package, install, and test your application on the application server:

- Install and uninstall applications
- Start and stop servers in a base configuration
- Run administrative scripts or commands
- Run the Enterprise JavaBeans (EJB) deployment tool for EJB 1.x or 2.x modules
- Run the JavaServer Pages (JSP) file precompilation tool

For more detailed information about Ant, refer to the Apache organization website.

Procedure

- To run Ant and have it automatically see the WebSphere classes, use the ws_ant command.
The ws_ant command is provided with the Apache Ant tool.
See the `app_server_root/bin/ws_ant.bat` file for the Apache Ant tool.
- Use Ant tasks for deployment and server operation.
The Apache Ant tasks for the product reside in the Java package: `com.ibm.websphere.ant.tasks`. The API documentation for this package contains detailed information about all of the Ant tasks that are provided and how to use them.
See `com.ibm.websphere.ant.tasks` API documentation in the **Reference** section of the information center.
- Use Ant tasks for building application code.

Refer to the Rational Application Developer documentation.

- Use the Apache Struts framework to create an extensible development environment for your application, based on published standards and proven design patterns.

Apache struts is a framework that is supported by the open source community.

IBM WebSphere Application Server provides the Apache Struts JAR file in the `install_root/optionalLibraries/Apache/Struts/1.1` directory of your product installation. If using Struts in your application or Application Server, you need to configure a shared library that points to the Struts library JAR file.

The Struts framework provides the invisible underpinnings every professional web application needs to survive. The core of Struts is a flexible control layer based on standard technologies such as Java Servlets, JavaBeans, ResourceBundles, and Extensible Markup Language (XML).

Note: Shipment of Apache Struts 1.1, 1.2.4, and 1.2.7 as optional libraries within WebSphere Application Server is deprecated in Version 7.0.

Struts encourages application architectures based on the Model 2 approach, a variation of the classic Model-View-Controller (MVC) design paradigm. Struts provides its own Controller component and integrates with other technologies to provide the Model and the View. For the Model, Struts can interact with any standard data access technology, including Enterprise Java Beans (EJB) components, and JDBC. For the View, Struts works well with JavaServer Pages (JSP) files, XSLT, or other presentation systems.

Chapter 5. Starting and stopping quick reference

Start and stop servers in your application serving environment, referring to this quick guide to the administrative clients and several other tools that are provided with this product.

Procedure

- Use commands to start and stop servers.

Table 13. Commands to start and stop servers. Run a start or stop command that is appropriate for the target server.

Quick reference: Issuing commands to start and stop servers
--

These examples are for starting and stopping the default profile on a server. Otherwise, you might need to specify <code>-profileName profile_name</code> when invoking the command.
--

Deployment manager

Run the following command.

See “startManager command” on page 115 for details and variations

<code>startManager</code>

Application server

Run the following command. See “startServer command” on page 110 for details and variations

<code>startServer server</code>

where <i>server</i> is the application server that you want to start.

Stopping the servers

Use the same command as to start, except substitute stop for start. For example, to stop an application server, issue the command:
--

<code>stopServer server</code>

To start and stop application server clusters, see Starting clusters.

- Use administrative clients and tools.

Table 14. Opening the administrative console. Point a Web browser at the console.

Quick reference: Opening the administrative console
--

To open the console using the default port, enter this web address in your web browser:

<code>http://your_fully_qualified_server_name:9060/ibm/console</code>

Depending on your configuration, your web address might differ. Other factors can affect your ability to access the console. See “Starting and logging off the administrative console” on page 20 for details, as needed.

- You also can use MVS console commands on z/OS systems.
 - To launch a scripting client, see Starting the wsadmin scripting client using wsadmin scripting.
 - To learn about all available administrative clients, see Using the administrative clients.
 - For performance monitoring, see Monitoring performance with Tivoli Performance Viewer .
- See the administrator commands that are listed in the **Reference** section of the information center.
- Use troubleshooting tools.
 - See Choosing and using diagnosis tools and controls on z/OS.

Chapter 6. Backing up and recovering the application serving environment

The product uses many operating system and application resources that you should consider adding to your backup and recovery procedures.

About this task

WebSphere Application Server resources can be saved while the product environment is active. When backing up database data, you may have to shut down some or all services if a snapshot cannot be obtained. This would occur if there are requests which obtain locks or have open transactions against the database being saved. In a distributed environment, you may need to consider how to get a consistent backup across several systems. If the data on systems is not closely related to data on other systems, you may be able to backup each system in isolation. If you need a snapshot across systems simultaneously, you may need to stop activity on all systems while the snapshot is taken.

How often you back up resources depends largely on when or how often you expect them to change.

Procedure

- Back up your product environment configuration.

This category covers the resources that define your WebSphere Application Server operating environment. Once you have done initial setup, this information should change very infrequently. You might backup this information only when you change these settings, and not include these resources in regularly scheduled backups.

- Administrative configuration files
- HTTP configuration (see the documentation for your web server)

- Back up your applications.

This category covers the applications you run using the product. You should back these up the same way you back up other applications on your system. You could backup these resources every time you add or change an application, or include these resources in a regularly scheduled backup.

- Application deployment configuration files

- Back up your application data.

This category covers the data stores used by your WebSphere Application Server applications. Unless your applications serve only static information, these resources are usually quite dynamic. You should back these up the same way you back up other business data on your system. These resources are suited for inclusion in a regularly scheduled backup.

- Servlet session data

What to do next

If your applications are using other resources or services that are external to the product, remember to include those in your backup plan as well.

Chapter 7. Class loading

Class loaders are part of the Java virtual machine (JVM) code and are responsible for finding and loading class files. Class loaders enable applications that are deployed on servers to access repositories of available classes and resources. Application developers and deployers must consider the location of class and resource files, and the class loaders used to access those files, to make the files available to deployed applications. Class loaders affect the packaging of applications and the runtime behavior of packaged applications of deployed applications.

Before you begin

This topic describes how to configure class loaders for application files or modules that are installed on an application server.

To better understand class loaders in WebSphere Application Server, read “Class loaders.” The topic “Class loading: Resources for learning” on page 203 refers to additional sources.

About this task

Configure class loaders for application files or modules that are installed on an application server using the administrative console. You configure class loaders to ensure that deployed application files and modules can access the classes and resources that they need to run successfully.

Procedure

1. If an installed application module uses a resource, create a resource provider that specifies the directory name of the resource drivers.
Do not specify the resource Java archive (JAR) file names. All JAR files in the specified directory are added into the class path of the WebSphere Application Server extensions class loader. If a resource driver requires a native library (.dll or .so file), specify the name of the directory that contains the library in the native path of the resource configuration.
2. Specify class-loader values for an application server.
3. Specify class-loader values for an installed enterprise application.
4. Specify the class-loader mode for an installed web module.
5. If your deployed application uses shared library files, associate the shared library files with your application. Use a library reference to associate a shared library file with your application.
 - a. If you have not done so already, define shared libraries for library files that your applications need.
 - b. Define a library reference for each shared library that your application uses.

What to do next

After configuring class loaders, ensure that your application performs as desired. To diagnose and fix problems with class loaders, refer to Troubleshooting class loaders.

Class loaders

Class loaders find and load class files. Class loaders enable applications that are deployed on servers to access repositories of available classes and resources. Application developers and deployers must consider the location of class and resource files, and the class loaders used to access those files, to make the files available to deployed applications.

This topic provides the following information about class loaders in WebSphere Application Server:

- “Class loaders used and the order of use” on page 194
- “Class-loader isolation policies” on page 195

- “Class-loader modes” on page 197

Class loaders used and the order of use

The product runtime environment 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 Java virtual machine

The bootstrap class loader uses the boot class path (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 loads the Java Platform, Enterprise Edition (Java EE) application programming interfaces (APIs) provided by the WebSphere Application Server product in the `j2ee.jar` file. Because this class loader loads the Java EE APIs, you can add libraries that depend on the Java EE APIs to the class path system property to extend a server class path. However, a preferred method of extending a server class path is to add a shared library.

2. A WebSphere extensions class loader

The WebSphere extensions class loader loads the WebSphere Application Server classes that are required at run time. The extensions class loader uses a `ws.ext.dirs` system property to determine the path that is used to load classes. Each directory in the `ws.ext.dirs` class path and every Java archive (JAR) file or compressed file in these directories is added to the class path 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, enterprise bean (EJB) modules, resource adapter archives (RAR files), and dependency JAR files. Application class loaders follow Java EE class-loading rules to load classes and JAR files from an enterprise application. The product enables you to associate shared libraries with an application.

4. Zero or more web module class loaders

By default, web module class loaders load the contents of the `WEB-INF/classes` and `WEB-INF/lib` directories. Web module class loaders are children of application class loaders. You can specify that an application class loader load the contents of a web module rather than the web module class loader.



Each class loader is a child of the previous class loader. That is, the application module class loaders are children of the WebSphere 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 extensions class loader is requested to find a class in a Java EE module, it cannot go to the application module class loader to find that class and a `ClassNotFoundException` error occurs. After 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 depend on the class-loader policies that are 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:

Table 15. Class-loader policy descriptions. Available policies include Application and WAR.

Class-loader policy	Description
Application	Application class loaders load EJB modules, dependency JAR files, embedded resource adapters, and application-scoped shared libraries. Depending on the application class-loader policy, an application class loader can be shared by multiple applications (<code>Single</code>) or unique for each application (<code>Multiple</code>). The application class-loader policy controls the isolation of applications that are running in the system. When set to <code>Single</code> , applications are not isolated. When set to <code>Multiple</code> , applications are isolated from each other.
WAR	<p>By default, web module class loaders load the contents of the <code>WEB-INF/classes</code> and <code>WEB-INF/lib</code> directories. The application class loader is the parent of the web module class loader. You can change the default behavior by changing the web application archive (WAR) class-loader policy of the application.</p> <p>The WAR class-loader policy controls the isolation of web modules. If this policy is set to <code>Application</code>, 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 <code>Module</code>, then each web module receives its own class loader whose parent is the application class loader.</p> <p>Tip: The console and the underlying <code>deployment.xml</code> file use different names for WAR class-loader policy values. In the console, the WAR class-loader policy values are <code>Application</code> or <code>Module</code>. However, in the underlying <code>deployment.xml</code> file where the policy is set, the WAR class-loader policy values are <code>Single</code> instead of <code>Application</code>, or <code>Multiple</code> instead of <code>Module</code>. <code>Application</code> is the same as <code>Single</code>, and <code>Module</code> is the same as <code>Multiple</code>.</p>

Restriction: WebSphere Application 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 that is used for loading the EJB modules, dependency JAR files, and shared libraries for that application.

An application class loader loads classes from web modules if the application's WAR class-loader policy is set to `Application`. If the application's 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 of the 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 that have a WAR class-loader policy set to `Module` use a separate class loader for web modules.

```
Server's application class-loader policy: Single
Application's WAR class-loader policy: Module
```

```
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
```

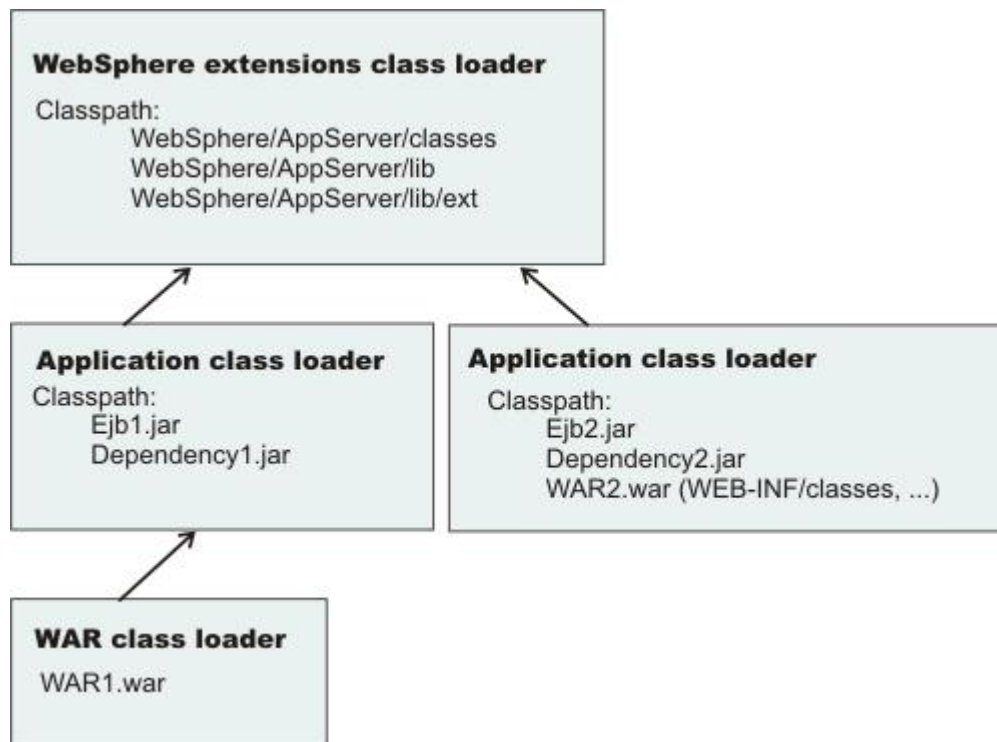


The following example shows that when the application class-loader policy of an application server is set to `Multiple`, each application on the server has its own class loader. An application class loader also loads its web modules if the application WAR class-loader policy is set to `Application`. If the policy is set to `Module`, then a web module uses its own class loader.

```
Server's application class-loader policy: Multiple
Application's WAR class-loader policy: Module
```

```
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

The class-loader delegation mode, also known as the *class loader order*, determines whether a class loader delegates the loading of classes to the parent class loader. The following values for class-loader mode are supported:

Table 16. Class-loader mode descriptions. Available modes include Parent first and Parent last.

Class-loader mode	Description
Parent first Also known as Classes loaded with parent class loader first.	The Parent first class-loader mode 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. This value is the default for the class-loader policy and for standard JVM class loaders.
Parent last Also known as Classes loaded with local class loader first or Application first.	The Parent last class-loader mode causes the class loader to attempt to load classes from its local class path before delegating the class loading to its parent. Using this policy, an application class loader can override and provide its own version of a class that exists in the parent class loader.

The following settings determine the mode of a class loader:

- If the application class-loader policy of an application server is `Single`, the server-level mode value defines the mode for an application class loader.
- If the application class-loader policy of an application server is `Multiple`, the application-level mode value defines the mode for an application class loader.

- If the WAR class-loader policy of an application is `Module`, the module-level mode value defines the mode for a WAR class loader.

Configuring class loaders of a server

You can configure the application class loaders for an application server. Class loaders enable applications that are deployed on the application server to access repositories of available classes and resources.

Before you begin

This topic assumes that an administrator created an application server on a WebSphere Application Server product.

About this task

Configure the class loaders of an application server to set class-loader policy and mode values which affect all applications that are deployed on the server. Use the administrative console to configure the class loaders.

Procedure

1. Click **Servers > Server Types > WebSphere application servers > *server_name*** to access an application server settings page.
2. Specify the application class-loader policy for the application server.

The application class-loader policy controls the isolation of applications that run in the system (on the server). An application class loader groups enterprise bean (EJB) modules, shared libraries, resource adapter archives (RAR files), and dependency Java archive (JAR) files associated to an application. Dependency JAR files are JAR files that contain code which can be used by both enterprise beans and servlets. The application class-loader policy controls whether an application class loader can be shared by multiple applications or is unique for each application.

Use the application server settings page to specify the application class-loader policy for the server:

Option	Description
Single	Applications are not isolated from each other. Uses a single application class loader to load all of the EJB modules, shared libraries, and dependency JAR files in the system.
Multiple	Applications are isolated from each other. Gives each application its own class loader to load the EJB modules, shared libraries, and dependency JAR files of that application.

3. Specify the application class-loader mode for the application server.
The application class loading mode specifies the class-loader mode when the application class-loader policy is `Single`.

On the application server settings page, select either of the following values:

Option	Description
Classes loaded with parent class loader 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. Classes loaded with parent class loader first is the default value for class loading mode. This value is also known as <code>parent first</code> .

Option	Description
Classes loaded with local class loader first (parent last)	Causes the class loader to attempt to load classes from its local class path before delegating the class loading to its parent. Using this policy, an application class loader can override and provide its own version of a class that exists in the parent class loader.

4. Specify the class-loader mode for the class loader.

- a. On the application server settings page, click **Java and Process Management > Class loader** to access the Class loader page.
- b. On the Class loader page, click **New** to access the settings page for a class loader.
- c. On the class loader settings page, specify the class loader order.

The `Classes loaded with parent class loader first` value 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.

The `Classes loaded with local class loader first (parent last)` value causes the class loader to attempt to load classes from its local class path before delegating the class loading to its parent.

- d. Click **OK**.

An identifier is assigned to a class-loader instance. The instance is added to the collection of class loaders shown on the Class loader page.

What to do next

Save the changes to the administrative configuration.

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 > Server types > WebSphere application servers > *server_name***. Under **Server Infrastructure**, expand **Java and Process Management** then click **Class loader**.

Class loader ID

Specifies a string that is unique to the server identifying the class-loader instance. The product assigns the identifier.

Class loader order

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 development kit class loaders and WebSphere Application Server class loaders is `Classes loaded with parent class loader first (Parent first)`. By specifying `Classes loaded with local class loader first (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.

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 > Server types > WebSphere application servers > *server_name***. Under **Server Infrastructure**, expand **Java and Process Management** then click **Class loader**. Click on a *Class_loader_ID*.

Class loader ID

Specifies a string that is unique to the server identifying the class-loader instance. The product assigns the identifier.

Data type String

Class loader order

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 development kit class loaders and WebSphere Application Server class loaders is *Classes loaded with parent class loader first*. By specifying *Classes loaded with local class loader first (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 *Classes loaded with parent class loader first* and *Classes loaded with local class loader first (parent last)*. The default is to search in the parent class loader before searching in the application class loader to load a class.

For your application to use the default configuration of Jakarta Commons Logging in this product, set this application class loader order to *Classes loaded with parent class loader first*. For your application to override the default configuration of Jakarta Commons Logging, your application must provide the configuration in a form supported by Jakarta Commons Logging and this class loader order must be set to *Classes loaded with local class loader first (parent last)*. Also, to override the default configuration, set the class loader order for each web module in your application so that the correct logger factory loads.

Data type String
Default Parent first

Configuring application class loaders

You can set values that control the class-loading behavior of an installed enterprise application. Class loaders enable an application to access repositories of available classes and resources.

Before you begin

This topic assumes that you installed an application on an application server.

About this task

Configure the class loaders of an enterprise application to set class-loader policy and mode values for this application.

An application class loader groups enterprise bean (EJB) modules, shared libraries, resource adapter archive (RAR) files, and dependency Java archive (JAR) files associated to an application. Dependency JAR files are JAR files that contain code which can be used by both enterprise beans and servlets.

An application class loader is the parent of a web application archive (WAR) class loader. By default, a web module has its own WAR class loader to load the contents of the web module. The WAR class-loader policy value of an application class loader determines whether the WAR class loader or the application class loader is used to load the contents of the Web module.

Use the administrative console to configure the class loaders.

Attention: If an application is running, changing an application setting causes the application to restart. On stand-alone servers, the application restarts after you save the change. On multiple-server products, the application restarts after you save the change and files synchronize on the node where the application is installed. To control when synchronization occurs on multiple-server products, deselect **Synchronize changes with nodes** on the Console preferences page.

Procedure

1. Click **Applications > Application Types > WebSphere enterprise applications > application_name > Class loading and update detection** to access the settings page for an application class loader.

2. Specify whether to reload application classes when the application or its files are updated.

By default, class reloading is not enabled. Select **Override class reloading settings for web and EJB modules** to choose to reload application classes. You might specify different values for EJB modules and for web modules such as servlets and JavaServer Pages (JSP) files.

3. Specify the number of seconds to scan the application's file system for updated files.

The value specified for **Polling interval for updated files** takes effect only if class reloading is enabled. The default is the value of the reloading interval attribute in the IBM extension (META-INF/ibm-application-ext.xmi) file of the enterprise application (EAR file). You might specify different values for EJB modules and for web modules such as servlets and JSP files.

To enable reloading, specify an integer value that is greater than zero (for example, 1 to 2147483647).

To disable reloading, specify zero (0).

Note: For IBM extension and binding files, the .xmi or .xml file name extension is different depending on whether you are using a pre-Java EE 5 application or module or a Java EE 5 or later application or module. An IBM extension or binding file is named ibm-*-ext.xmi or ibm-*-bnd.xmi where * is the type of extension or binding file such as app, application, ejb-jar, or web. The following conditions apply:

- For an application or module that uses a Java EE version prior to version 5, the file extension must be .xmi.
- For an application or module that uses Java EE 5 or later, the file extension must be .xml. If .xmi files are included with the application or module, the product ignores the .xmi files.

However, a Java EE 5 or later module can exist within an application that includes pre-Java EE 5 files and uses the .xmi file name extension.

The ibm-webservices-ext.xmi, ibm-webservices-bnd.xmi, ibm-webservicesclient-bnd.xmi, ibm-webservicesclient-ext.xmi, and ibm-portlet-ext.xmi files continue to use the .xmi file extensions.

4. Specify the class loader order for the application.

The application class loader order specifies whether the class loader searches in the parent class loader or in the application class loader first to load a class. The default is to search in the parent class loader before searching in the application class loader to load a class.

Select either of the following values for **Classes loader order**:

Option	Description
Classes loaded with parent class loader first	Causes the class loader to search in the parent class loader first to load a class. This value is the standard for Development Kit class loaders and WebSphere Application Server class loaders.

Option	Description
Classes loaded with local class loader first (parent last)	Causes the class loader to search in the application class loader first to load a class. By specifying <code>Classes loaded with local class loader first (parent last)</code> , your application can override classes contained in the parent class loader. Note: Specifying the <code>Classes loaded with local class loader first (parent last)</code> value might result in <code>LinkageErrors</code> or <code>ClassCastException</code> messages if you have mixed use of overridden classes and non-overridden classes.

- Specify whether to use a single or multiple class loaders to load web application archives (WAR files) of your application.

By default, web modules have their own WAR class loader to load the contents of the `WEB-INF/classes` and `WEB-INF/lib` directories. The default WAR class loader value is `Class loader for each WAR file in application`, which uses a separate class loader to load each WAR file. Setting the value to `Single class loader for application` causes the application class loader to load the web module contents as well as the EJB modules, shared libraries, RAR files, and dependency JAR files associated to the application. The application class loader is the parent of the WAR class loader.

Select either of the following values for **WAR class loader policy**:

Option	Description
Class loader for each WAR file in application	Uses a different class loader for each WAR file.
Single class loader for application	Uses a single class loader to load all of the WAR files in your application.

- Click **OK**.

What to do next

Save the changes to the administrative configuration.

Configuring web module class loaders

You can set values that control the class-loading behavior of an installed web module.

Before you begin

This topic assumes that you installed a web module on an application server.

About this task

Configure the class loader order value of an installed web module. By default, a web module has its own web application archive (WAR) class loader to load the contents of the web module, which are in the `WEB-INF/classes` and `WEB-INF/lib` directories.

An application class loader is the parent of a WAR class loader. The WAR class-loader policy value of an application class loader determines whether the WAR class loader or the application class loader is used to load the contents of the web module.

The default WAR class loader policy value is `Class loader for each WAR file in application`. If the policy is set to `Class loader for each WAR file in application`, then each web module receives its own class loader whose parent is the application class loader. If the policy is set to `Single class loader for application`, then the application class loader loads the web module contents as well as the enterprise

bean (EJB) modules, shared libraries, resource adapter archive (RAR) files, and dependency Java archive (JAR) files associated to an application. Thus, the configuration of the parent application class loader affects the WAR class loader. You can set the policy on the Class loading and update detection page of an administrative console.

Use the administrative console to configure the application and WAR class loaders.

Attention: If an application is running, changing an application setting causes the application to restart. On stand-alone servers, the application restarts after you save the change. On multiple-server products, the application restarts after you save the change and files synchronize on the node where the application is installed. To control when synchronization occurs on multiple-server products, deselect **Synchronize changes with nodes** on the Console preferences page.

Procedure

1. If you have not done so already, configure the application class loader.

Settings such as **Override class reloading settings for web and EJB modules**, **Polling interval for updated files** and **WAR class loader policy** can affect web module class loading.

If **WAR class loader policy** is set to **Class loader** for each WAR file in application, then the web module receives its own class loader and the WAR class-loader policy of the web module defines the mode for a WAR class loader. If the policy is set to **Single class loader** for application, then the application class loader loads the web module contents.

2. Specify the class loader order for the installed web module.

The web module class-loader mode specifies whether the class loader searches in the parent application class loader or in the WAR class loader first to load a class. The default is to search in the parent application class loader before searching in the WAR class loader to load a class.

Select either of the following values for **Class loader order**:

Option	Description
Classes loaded with parent class loader first	This option causes the class loader to <i>prefer classes that are provided by the product</i> over the classes that exist within the web module. This approach is standard for Development Kit class loaders and WebSphere Application Server class loaders.
Classes loaded with local class loader first	This option causes the class loader to <i>prefer classes that exist in the web module</i> over the classes that are provided by the product. If the same class exists in both the product and the web module, the class from the web module is loaded. Attention: If you specify the Classes loaded with local class loader first value, you might receive <code>LinkageErrors</code> or <code>ClassCastException</code> messages if you have mixed use of overridden classes and non-overridden classes.

3. Click **OK**.

What to do next

Save the changes to the administrative configuration.

Class loading: Resources for learning

Additional information and guidance on class loading is available on various Internet sites.

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.

View links to additional information about:

- “Programming model and decisions”
- “Programming instructions and examples”
- “Programming specifications”

Programming model and decisions

- Demystifying class loading problems, Part 1: An introduction to class loading and debugging tools - Learn how class loading works and how your JVM can help you sort out class loading problems (*developerWorks*, November 2005), http://www.ibm.com/developerworks/java/library/j-dclp1/?S_TACT=106AH10W&S_CMP=NC
- Demystifying class loading problems, Part 2: Basic class loading exceptions - An in-depth look at some simple class loading quirks and conundrums (*developerWorks*, December 2005), http://www.ibm.com/developerworks/java/library/j-dclp2.html?S_TACT=105AGX10&S_CMP=NC
- Demystifying class loading problems, Part 3: Tackling more unusual class loading problems - Understand class loading and quash subtle exceptions (*developerWorks*, December 2005), http://www.ibm.com/developerworks/java/library/j-dclp3/?S_TACT=105AGX10&S_CMP=NC
- J2EE Class Loading Demystified (*developerWorks*, August 2002), http://www.ibm.com/developerworks/websphere/library/techarticles/0112_deboer/deboer.html
- Java programming dynamics, Part 1: Classes and class loading - A look at classes and what goes on as they're loaded by a JVM (*developerWorks*, April 2003), <http://www.ibm.com/developerworks/java/library/j-dyn0429/>

Programming instructions and examples

- *WebSphere Application Server V6.1: System Management Configuration Handbook*, SG24-7304-00, <http://www.redbooks.ibm.com/abstracts/SG247304.html?open>
- *IBM WebSphere Developer Technical Journal: Co-hosting multiple versions of J2EE applications*, http://www.ibm.com/developerworks/websphere/techjournal/0405_poddar/0405_poddar.html

Programming specifications

- Specifications and API documentation

Chapter 8. Deploying and administering enterprise applications

Deploying an enterprise application file consists of installing an application file on a server configured to hold installable Java Platform, Enterprise Edition (Java EE) modules.

Before you begin

Before installing an enterprise application or other installable module on an application server, you must develop the module, assemble the module, and configure the target server or cluster. Before choosing a deployment target for the module, ensure that the target version is compatible with your module.

About this task

During installation, you can configure the module enough to enable it to run on the server. After installation, you can configure the module further, start or stop the application, and otherwise manage its activity.

The topics in this section describe how to deploy and administer applications or modules using the administrative console. You can also use scripting or administrative programs (JMX).

Procedure

- Install Java EE application files on an application server.
- Edit the administrative configuration for an application.
- Optional: View the deployment descriptor for an application or module.
- Start and stop enterprise applications.
- Export enterprise applications.
- Export a file in a Java EE application or module.
- Export DDL files.
- Update a Java EE application or module.
- Uninstall enterprise applications using the console.
- Uninstall enterprise applications by removing them from a monitored directory.
- Remove a file from a Java EE application or module.

What to do next

If a changed application or module is deployed on a cluster, click **Rollout Update** on the Enterprise applications page to propagate the changed configuration on all cluster members of the cluster on which the application or module is deployed. **Rollout Update** sequentially updates the configuration on the nodes that contain cluster members.

Use **Rollout Update** if the application is deployed on one or more clusters spread across multiple nodes. This action reduces the amount of time that any single cluster member is unavailable to serve requests to the smallest interval possible. Pending I/O transactions will complete before a cluster member stops; in-flight HTTP and Java Message Service (JMS) transactions might be lost while the cluster member is stopping. If the administrative console session times out during a rollout updating, ensure that the changed configuration was propagated to each node. For an application server without clusters, use **Update** and then save and synchronize the node instead. For a stand-alone application server, simply update and save.

After making changes to administrative configurations of your applications in the administrative console, ensure that you save the changes.

Enterprise (Java EE) applications

Enterprise applications (or Java EE applications) are applications that conform to the Java Platform, Enterprise Edition (Java EE) specification. Prior to Java EE 5, the specification name was Java 2 Platform, Enterprise Edition (J2EE). The term *Java EE* includes Java EE 5 and J2EE specifications.

Enterprise applications can consist of the following:

- Zero or more EJB modules (packaged in JAR files)
- Zero or more web modules (packaged in WAR files)
- Zero or more connector modules (packaged in RAR files)
- Zero or more Session Initiation Protocol (SIP) modules (packaged in SAR 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 Java EE application is represented by, and packaged in, an enterprise archive (EAR) file.

System applications

A *system application* is a Java Platform, Enterprise Edition (Java EE) enterprise application that is central to a WebSphere Application Server product.

Examples of system applications include *isclite*, *managementEJB* and *filetransfer*.

Because a system application is an important part of a WebSphere Application Server product, a system application is deployed when the product is installed and is updated only through a product fix or upgrade. For some system applications, such as *filetransfer*, users cannot change the metadata for the system application, unless the metadata assigns users and groups for security purposes. For these applications, non-security related metadata such as its Java EE bindings or extensions must be updated through a product fix or upgrade.

System applications are not shown in the list of installed applications on the console Enterprise Applications page, or through wsadmin and Java application programming interfaces, to prevent users from accidentally stopping, updating or removing the system applications.

Note that Java EE Samples are not system applications even though they are provided as part of a WebSphere Application Server product. Similarly, applications that support changes to their metadata are not system applications.

Common deployment framework

The *common deployment framework* enables you to implement plug-ins that add steps to default Java Platform, Enterprise Edition (Java EE) application management operations such as install, uninstall, edit and update.

Using the framework, you can implement management operations on specific types of deployable contents. For example, the deployable contents might include EAR, WAR, JAR or other Java EE modules and the management operations might include install and uninstall. Each operation is divided into a number of steps. For example, the install operation has steps for EJBDeploy and JavaServer Pages (JSP) compilation, among others. Using the common deployment framework, you can add steps to the default logic for Java EE operations.

The product supports framework plug-ins that extend deployment of EAR files. An EAR file has operations such as createEarWrapper, installApplication, uninstallApplication and editApplication. Using a framework plug-in, you can add steps to default install operations that support, for example, creating additional configuration artifacts in a configuration session, modifying an input EAR file using code generation, or additional validating of input parameters.

To extend application management operations using the framework, a plug-in must do the following:

- Implement each step.

A *step* runs logic that performs an operation. A step can access the deployment context and the deployable object. The *deployment context* provides information such as the operation name, the configuration session identifier, the temporary location for creating temporary files, operations parameters, and the like. A step is added by the extension provider.

- Implement an extension provider that adds each implemented step.

An *extension provider* is a class that provides steps for an operation on a given type, the EAR file type.

- Register the plug-in with a WebSphere Application Server server.

The plug-in is implemented as an Eclipse plug-in and is placed in *app_server_root/plugins* directory.

Add the extension point for the extension provider in the META-INF/plugin.xml file within the plug-in JAR file.

For an example of these steps, refer to Extending application management operations through programming.

Installing enterprise application files

As part of deploying an application, you install application files on a server configured to hold installable modules.

Before you begin

Before you can install your Java Platform, Enterprise Edition (Java EE) application files on an application server, you must assemble modules as needed.

Also, before you install the files, configure the target application server. As part of configuring the server, determine whether your application files can be installed to your deployment targets.

Note: Installation of new Java EE specification levels are not allowed on deployment targets that do not support the specified Java EE level. If an application contains modules with an EJB deployment descriptor level of 3.1 (part of Java EE 6), the application cannot be installed on an application server that does not support Java EE 6; for example, application server versions prior to Version 8. Further, if an application contains Java EE 6 annotations, the application cannot be installed on an application server that does not support the Java EE 6 specification level, even if the application contains an EJB deployment descriptor for a previous version of the supported Java EE specification level; for example, EJB 3.0 which is part of Java EE 5.

About this task

You can install the following enterprise modules on a server:

- Enterprise archive (EAR)
- Enterprise bean (EJB)
- Web archive (WAR)
- Session Initiation Protocol (SIP) module (SAR)
- Resource adapter (connector or RAR)
- Application client modules

Application client files can be installed in a WebSphere Application Server configuration but cannot be run on a server.

Complete the following steps to install your files.

Procedure

1. Determine which method to use to install your application files. The product provides several ways to install modules.
2. Install the application files using
 - Administrative console
 - Drag and drop to a monitored directory
 - wsadmin scripts
 - Java administrative programs that use Java Management Extensions (JMX) application programming interfaces (APIs)
 - Java programs that define a Java EE DeploymentManager object in accordance with Java EE Application Deployment specification (JSR-88)
3. Start the deployed application files using
 - Administrative console
 - wsadmin startApplication
 - Java programs that use ApplicationManager or AppManagement MBeans
 - Java programs that define a Java EE DeploymentManager object in accordance with Java EE Application Deployment specification (JSR-88)

What to do next

Save the changes to your administrative configuration.

When saving the configuration in multiple-server environments, synchronize the configuration with the nodes where the application is expected to run.

Next, test the application. For example, point a web browser at the URL for a deployed application. Typically, the URL is `http://hostname:9060/web_module_name`, where *hostname* is your valid web server and 9060 is the default port number. Examine the performance of the application. If the application does not perform as desired, edit the application configuration, then save and test it again.

If your application contains many classes with annotations and takes a long time to deploy, you can reduce annotation searches to speed up deployment. See the topic on reducing annotation searches during application deployment.

Installable enterprise module versions

The contents of a Java Platform, Enterprise Edition (Java EE) module affect whether you can install the module on a deployment target. A *deployment target* is a server or a cluster with at least one member on a WebSphere Application Server product.

Installable application modules

Select only appropriate deployment targets for a module. You must install an application, enterprise bean (EJB) module, Session Initiation Protocol (SIP) archive (SAR), web module, or client module on a Version 8.x target under any of the following conditions:

- The module supports Java Platform, Enterprise Edition (Java EE) 6
- The module calls an 8.x runtime application programming interface (API).
- The module uses an 8.x product feature.

For example, because support for deployment of application client modules using the administrative console or wsadmin AdminApp commands was added in Version 8.0, you must install a client module using the console or an AdminApp command only to a Version 8.x target.

If a module supports Java 2 Platform, Enterprise Edition (J2EE) 1.4, then you can install the module on a Version 6.x, 7.x or 8.x deployment target. Modules that call a 6.1.x API or use a 6.1.x feature can be

installed on a 6.1.x, 7.x or 8.x deployment target. Modules that call a 6.0.x API or use a 6.0.x feature can be installed on a 6.0.x, 6.1.x, 7.x or 8.x deployment target. Modules that require 6.1.x feature pack functionality can be installed on a 7.x or 8.x deployment target or on a 6.1.x deployment target that has been enabled with that feature pack. Modules that require 7.x feature pack functionality can be installed on a 8.x deployment target or on a 7.x deployment target that has been enabled with that feature pack.

Selecting options such as **Precompile JavaServer Pages files**, **Use binary configuration**, **Deploy web services** or **Deploy enterprise beans** during application installation indicates that the application uses 6.1.x product features. You cannot deploy such applications on a 6.0.x deployment target. You must deploy such applications on a 6.1.x, 7.x or 8.x deployment target.

Note: You must package container-managed persistence (CMP) or bean-managed persistence (BMP) entity beans in an EJB 2.1 or earlier module. You cannot install an EJB 3.0 or EJB 3.1 module that contains CMP or BMP entity beans. Installation fails when a CMP or BMP entity bean is packaged in an EJB 3.0 or EJB 3.1 module. You can install EJB 2.1 or earlier modules on a 6.x, 7.x or 8.x deployment target.

Installable RAR files

You can install a stand-alone resource adapter (connector) module, or RAR file, developed for a Version 6.0.x product to a 6.x, 7.x or 8.x deployment target. If the module calls a 6.1.x API, then you must install the module on a 6.1.x, 7.x or 8.x deployment target. You must install a module that calls a 7.x API on a 7.x or 8.x deployment target. You must install a module that calls a 8.x API on a 8.x deployment target.

Deployment targets

Table 1 lists the compatible deployment target versions for various modules. "6.x, 7.x or 8.x" for **Deployment target versions** indicates that you can deploy the module to a WebSphere Application Server Version 6, 7, or 8 server or cluster.

Table 17. Compatible deployment target versions for 6.x, 7.x and 8.x modules. Deploy modules to compatible deployment target versions.

Module type	Module Java support	Module calls 6.x, 7.x or 8.x runtime APIs or uses 6.x, 7.x or 8.x features?	Client versions that can install module	Deployment target versions
Application, EJB, or web	J2EE 1.3	No	6.x, 7.x or 8.x	6.x, 7.x or 8.x
Application, EJB, or web	J2EE 1.3	Yes	6.x, 7.x or 8.x for 6.x APIs or features 7.x or 8.x for 7.x APIs or features 8.x for 8.x APIs or features	6.x, 7.x or 8.x You must install modules that call 6.1.x runtime APIs or use 6.1.x features on a 6.1.x, 7.x or 8.x deployment target. You can install modules that call 6.0.x runtime APIs or use 6.0.x features on any 6.x, 7.x or 8.x deployment target.
Application, EJB, SAR, or web	J2EE 1.4	Yes or No	6.x, 7.x or 8.x	6.x, 7.x or 8.x
Application, EJB, SAR, or web	Java EE 5	Yes or No	7.x or 8.x	7.x or 8.x

Table 17. Compatible deployment target versions for 6.x, 7.x and 8.x modules (continued). Deploy modules to compatible deployment target versions.

Application, EJB, SAR, or web	Java EE 6	Yes or No	8.x	8.x
Client	Any Java EE version	Yes or No	8.x	8.x
Resource adapter	JCA 1.0	No	6.x, 7.x or 8.x	6.x, 7.x or 8.x
Resource adapter	JCA 1.0	Yes	6.x, 7.x or 8.x	6.x, 7.x or 8.x You must install modules that call 6.1.x runtime APIs on a 6.1.x, 7.x or 8.x deployment target. You can install modules that call 6.0.x runtime APIs on any 6.x, 7.x or 8.x deployment target.
Resource adapter	JCA 1.5	Yes or No	6.x, 7.x or 8.x	6.x, 7.x or 8.x You must install modules that call 6.1.x runtime APIs on a 6.1.x, 7.x or 8.x deployment target. You can install modules that call 6.0.x runtime APIs on any 6.x, 7.x or 8.x deployment target.
Resource adapter	JCA 1.6	Yes or No	JCA 1.6 resource adapters can only be installed on 8.x. Resource adapter archive annotations are not supported on previous WebSphere Application Server releases.	JCA 1.6 resource adapters can only be installed on 8.x. Resource adapter archive annotations are not supported on previous WebSphere Application Server releases.

Ways to install enterprise applications or modules

The product provides several ways to install Java Platform, Enterprise Edition (Java EE) application files.

Installable files include enterprise archive (EAR), enterprise bean (EJB), web application archive (WAR), Session Initiation Protocol (SIP) archive (SAR), resource adapter (connector or RAR), and application client modules. They can be installed on a server or cluster. Application client files can be installed in a WebSphere Application Server configuration but cannot be run on a server.

Table 18. Ways to install application files. Deploy an application or module using the administrative console, wsadmin, programming, or deployment tools.

Option	Method	Modules	Comments	Starting after install
Administrative console install wizard See topics on installing enterprise application files with the console.	Click Applications > New application > New Enterprise Application in the console navigation tree and follow instructions in the wizard.	Files for all of the following modules: <ul style="list-style-type: none"> • EAR • EJB • WAR • SAR • RAR • Application client 	Provides one of the easier ways to install application files. For applications that do not require changes to the default bindings, after you specify the application file, expand Choose to generate default bindings and mappings , select Generate default bindings , click the Summary step, and then click Finish .	Click Start on the Enterprise applications page accessed by clicking Applications > Application Types > WebSphere enterprise applications in the console navigation tree.
Monitored directory	Add an EAR file or module to a dragDropDeployableApps subdirectory of an application server or deployment manager profile. For base (stand-alone) application servers, the monitored directory is the dragDropDeployableApps /servers/server_name directory of the application server profile. For deployment managers, a deployment manager profile has the following monitored directories: <ul style="list-style-type: none"> • dragDropDeployableApps /servers/server_name • dragDropDeployableApps /nodes/node_name /servers/server_name • dragDropDeployableApps /clusters/cluster_name 	Files for all of the following modules: <ul style="list-style-type: none"> • EAR • EJB • WAR • SAR Use this option for drag and drop deployment of Java EE 5.0 and later modules.	Provides one of the easier ways to install applications. You cannot specify bindings during deployment. For applications that require changes to the bindings, install the application using the administrative console install wizard, application properties files, or wsadmin scripts. Monitored directory deployment differs from rapid deployment tools in several ways: <ul style="list-style-type: none"> • Monitored directory deployment supports deployment to base and network deployment environments. Rapid deployment tools support deployment only to the base environment. • Monitored directory deployment does not start a new daemon. Rapid deployment tools start a separate process. • Monitored directory deployment supports deployment of Java EE 5 and later modules. Rapid deployment tools support assembly of J2EE 1.3 and 1.4 modules, and deployment of all Java EE module versions. • Monitored directory deployment supports use of a properties file to specify deployment options. Rapid deployment tools does not support use of a properties file. 	After application or module installation or update, the product starts the application or module automatically. For installations on a deployment manager monitored directory, the product does not start an application or module automatically if the node agent or server is not running.
Application properties files	Create a properties file that specifies to install or update application files. Then, run the wsadmin applyConfigProperties command in any of the following ways: <ul style="list-style-type: none"> • In a script • At a command prompt • By adding the properties file to a dragDropDeployableApps /deploymentProperties monitored directory 	Files for all of the following modules: <ul style="list-style-type: none"> • EAR • EJB • WAR • SAR 	For information on deploying applications using properties files, see the following topics: <ul style="list-style-type: none"> • Using application properties files to install, update, and delete enterprise application files • Installing enterprise application files by adding properties files to a monitored directory 	After application or module installation or update, the product starts the application or module automatically. For installations on a deployment manager monitored directory, the product does not start an application or module automatically if the node agent or server is not running.

Table 18. Ways to install application files (continued). Deploy an application or module using the administrative console, wsadmin, programming, or deployment tools.

Option	Method	Modules	Comments	Starting after install
wsadmin scripts	Invoke AdminApp object install commands in a script or at a command prompt.	Files for all of the following modules: <ul style="list-style-type: none"> • EAR • EJB • WAR • SAR • RAR • Application client 	"Getting started with scripting" in the Using the administrative clients PDF provides an overview of wsadmin.	<ul style="list-style-type: none"> • Invoke the AdminApp startApplication command. • Invoke the startApplication method on an ApplicationManager MBean using AdminControl.
Job manager runs wsadmin scripts	Invoke AdminTask.submitJob -jobType installApplication command in a script or at a command prompt.	Files for all of the following modules: <ul style="list-style-type: none"> • EAR • EJB • WAR • SAR • RAR • Application client 	Use the job manager to submit a job that installs the application. You can also submit jobs that start and stop the application at designated times. "Administering jobs in a flexible management environment using scripting" in the Using the administrative clients PDF provides an overview of the job manager.	<ul style="list-style-type: none"> • Invoke the AdminTask.submitJob -jobType startApplication command. • Invoke the AdminApp startApplication command. • Invoke the startApplication method on an ApplicationManager MBean using AdminControl.
Java application programming interfaces	Install programs by completing the steps in Installing an application through programming.	All EAR files	Use Java Management Extensions (JMX) MBeans to install the application. For an overview of Java MBean programming, see Managing applications through programming.	Start the application by calling the startApplication method on a proxy.
Rapid deployment tools Refer to topics under Rapid deployment of J2EE applications .	Briefly, do the following: <ol style="list-style-type: none"> 1. Update your J2EE application files. 2. Set up the rapid deployment environment. 3. Create a free-form project. 4. Launch a rapid deployment session. 5. Drop your updated application files into the free-form project. 	<p>J2EE modules at the J2EE 1.3 or 1.4 specification levels, including EAR files and the following stand-alone modules:</p> <ul style="list-style-type: none"> • EJB • WAR • SAR • RAR • Application client <p>The rapid deployment tools do not support the J2EE 1.2 or Java EE 5.0 and later specification levels. Use this option for drag and drop deployment of J2EE 1.3 or 1.4 modules. Unlike the monitored directory option, the rapid deployment tools do not support drag and drop deployment of Java EE 5.0 and later modules.</p>	<p>Rapid deployment tools offer the following advantages:</p> <ul style="list-style-type: none"> • You do not need to assemble your J2EE application files prior to deployment. • You do not need to use other installation tools mentioned in this table to deploy the files. <p>For a list of ways in which the rapid deployment tools differ from monitored directory deployment, see the monitored directory description in this table.</p>	Use any of the options in this table to start the application. Clicking Start on the Enterprise applications page is the easiest option.
Java programs	Code programs that use Java EE DeploymentManager (JSR-88) methods. Note: Application installation using JSR-88 has been deprecated in WebSphere Application Server Version 8.0. Use another way listed in this table to deploy applications or modules.	All Java EE modules, including EAR files and the following stand-alone modules: <ul style="list-style-type: none"> • EJB • WAR • SAR • RAR • Application client 	<ul style="list-style-type: none"> • Uses Java EE Application Deployment Specification (JSR-88). • Can customize modules using DConfigBeans. 	Call the Java EE DeploymentManager (JSR-88) start method in a program to start the deployed modules when the module's running environment initializes.

Installing enterprise application files with the console

Installing Java Platform, Enterprise Edition (Java EE) application files consists of placing assembled enterprise application, Web, enterprise bean (EJB), or other installable modules on a server or cluster configured to hold the files. Installed files that start and run properly are considered *deployed*.

Before you begin

Before installing enterprise application files, ensure that you are installing your application files onto a compatible deployment target. If the deployment target is not compatible, select a different target.

Optionally, determine whether the application that you are installing uses library files that other deployed applications also use. You can define a shared library for each of these shared files. Using shared libraries reduces the number of library file copies on your workstation or server.

About this task

To install new enterprise application files to a WebSphere Application Server configuration, you can use the following options:

- Administrative console
- wsadmin scripts
- Monitored directory deployment
- Application properties files
- Java MBean programs
- Java programs that call Java EE DeploymentManager (JSR-88) methods

This topic describes how to use the administrative console to install an application, EJB component, Session Initiation Protocol (SIP) archive (SAR), or web module.

Note: After you start completing steps in the application installation wizard, 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.

Procedure

1. Click **Applications > New application > New Enterprise Application** in the console navigation tree.
2. On the first Preparing for application installation page:
 - a. Specify the full path name of the source enterprise 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 web archive (WAR), SAR, or Java archive (JAR) file for installation.
 - b. Click **Next**.
3. On the second Preparing for application installation page:
 - a. Select whether to view all installation options.

Fast Path - Prompt only when additional information is required
Displays the module mapping step as well as any steps that require you to specify needed information to install the application successfully.

Detailed - Show all installation options and parameters
Displays all installation options.
 - b. Select whether to generate default bindings.

Select **Generate default bindings** to have the product supply default values for incomplete Java Naming and Directory (JNDI) and other application bindings. The product does not change existing bindings.

You do not need to specify JNDI values for EJB bean, local home, remote home, or business interfaces of EJB 3.x modules. The product assigns container default values during run time. Similarly, for any EJB reference within an EJB 3.x or a Web 2.4 or later module, you do not need to specify JNDI values because the product resolves the targets automatically during run time. Even when you select **Generate default bindings**, the product does not generate default values for those JNDI values but it does generate default values for other bindings such as virtual host.

You can customize default values used in generating default bindings. “Preparing for application installation binding settings” on page 223 describes available customization and provides sample bindings.

- c. Click **Next**. If security warnings are displayed, click **Continue**. The Install New Application pages are displayed. If you chose to generate default bindings, you can proceed to the Summary step. “Example: Installing an EAR file using the default bindings” on page 220 provides sample steps.

4. Specify values for installation options as needed.

You can click on a step number to move directly to that page instead of clicking **Next**. The contents of the application or module that you are installing determines which pages are available.

Table 19. Wizard page descriptions. The table describes each wizard page.

Page	Description
Select installation options	On the Select installation options page, provide values for the settings specific to the product. Default values are used if you do not specify a value.
Map modules to servers	On the Map modules to servers page, specify deployment targets where you want to install the modules contained in your application. Modules can be installed on the same deployment target or dispersed among several deployment targets. Each module must be mapped to a target server. On single-server products, a deployment target can be an application server or web server. On multiple-server products, a deployment target can be an application server, cluster of application servers or web server.
Provide options to compile JSPs	If the Precompile JavaServer Pages files setting is enabled on the Select installation options page and your application uses JavaServer Pages (JSP) files, then you can specify JSP compiler options on the Provide options to compile JSPs page.
Provide JNDI names for beans	On the Provide JNDI names for beans page, specify a JNDI name for each enterprise bean in every EJB 2.1 and earlier 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. As to EJB 3.x modules, you can specify JNDI names, local home JNDI names, remote home JNDI names, or no JNDI names. If you do not specify a value, the product provides a default value.
Bind EJB business	On the Bind EJB business page, you can specify business interface JNDI names for EJB 3.x modules. If you specified a JNDI name for a bean on the Provide JNDI names for beans page, do not specify a business interface JNDI name on this page for the same bean. If you do not specify the JNDI name for a bean, you can optionally specify a business interface JNDI name. When you do not specify a business interface JNDI name, the product provides a container default. For a no-interface view, the business interface value is an empty string ("").

Table 19. Wizard page descriptions (continued). The table describes each wizard page.

Page	Description
Map default data sources for modules containing 1.x entity beans	If your application uses EJB modules that contain Container Managed Persistence (CMP) beans that are based on the EJB 1.x specification, for Map default data sources 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.
Map EJB references to beans	<p>On the Map EJB references to beans page, if your application defines EJB references, you can specify JNDI names for enterprise beans that represent the logical names specified in EJB references.</p> <p>If the EJB reference is from an EJB 3.x, or Web 2.4 or later module, the JNDI name is optional. For earlier modules, each EJB reference defined in the application must be bound to an EJB file.</p> <p>If Allow EJB reference targets to resolve automatically is enabled, the JNDI name is optional for all modules. The product provides a container default value or automatically resolves the EJB reference for incomplete bindings.</p>
Map resource references to resources	If your application defines resource references, for Map resource references to resources , specify JNDI names for the resources that represent the logical names defined in resource references. You can optionally specify login configuration name and authentication properties for the resource. After specifying authentication properties, click OK to save the values and return to the mapping step. You can optionally specify extended data source properties to enable a data source that uses heterogeneous pooling to connect to a DB2® database. Each resource reference defined in the application must be bound to a resource defined in your WebSphere Application Server configuration before clicking Finish on the Summary page.
Map virtual hosts for web modules	If your application uses web modules, for Map virtual hosts for web modules , select a virtual host from the list to 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 page.
Map security roles to users or groups	<p>If the application has security roles defined in its deployment descriptor then, for Map security roles to users or groups, specify users and groups that are mapped to each of the security roles. Select Role to select all the roles or select individual roles. For each role, you can specify whether predefined users such as Everyone or All authenticated users are mapped to it. To select specific users or groups from the user registry:</p> <ol style="list-style-type: none"> 1. Select a role and click Lookup users or Lookup groups. 2. On the Lookup users or groups page displayed, enter search criteria to extract a list of users or groups from the user registry. 3. Select individual users or groups from the results displayed. 4. Click OK to map the selected users or groups to the role selected on the Map security roles to users or groups page.
Map RunAs roles to users	If the application has Run As roles defined in its deployment descriptor, for Map RunAs roles to users , 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 Role to select all the roles or select individual roles. After selecting a role, enter values for the user name, password, and verify password and click Apply .
Ensure all unprotected 1.x methods have the correct level of protection	If your application contains EJB 1.x CMP beans that do not have method permissions defined for some of the EJB methods, for 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.

Table 19. Wizard page descriptions (continued). The table describes each wizard page.

Page	Description
Bind listeners for message-driven beans	If your application contains message driven enterprise beans, for Bind listeners for message-driven beans , provide a listener port name or an activation specification JNDI name for every message driven bean.
Map default data sources for modules containing 2.x entity beans	If your application uses EJB modules that contain CMP beans that are based on the EJB 2.x specification, for Map default data sources for modules containing 2.x 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. You can optionally specify a login configuration name and authentication properties for the data source. When creating authentication properties, you must click OK to save the values and return to the mapping step. You can optionally specify extended data source properties to enable a data source that uses heterogeneous pooling to connect to a DB2 database. The default data source for EJB modules is optional if data sources are specified for individual CMP beans.
Map data sources for all 2.x CMP beans	If your application has CMP beans that are based on the EJB 2.x specification, on the Map data sources for all 2.x CMP beans page, for each of the 2.x CMP beans specify a JNDI name and the type of resource authorization for data sources to be used. You can optionally specify a login configuration name and authentication properties for the data source. When creating authentication properties, you must click OK to save the values and return to the mapping step. 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 a default data source for the EJB module and a data source for individual CMP beans are not specified, then a validation error is displayed after you click Finish and installation is canceled.
Ensure all unprotected 2.x methods have the correct level of protection	If your application contains EJB 2.x CMP beans that do not have method permissions defined in the deployment descriptors for some of the EJB methods, on the Ensure all unprotected 2.x methods have the correct level of protection page, specify whether you want to assign a specific role to the unprotected methods, add the methods to the exclude list, or mark them as deselected. Methods added to the exclude list are marked as uncallable. For methods marked deselected no authorization check is performed before their invocation.
Provide options to perform the EJB Deploy	If the Deploy enterprise beans setting is enabled on the Select installation options page, then you can specify options for the EJB deployment tool on the Provide options to perform the EJB Deploy page. On this page, you can specify extra class paths, RMIC options, database types, and database schema names to be used while running the EJB deployment tool. You can specify the EJB deployment tool options on this page when installing or updating an application that contains EJB modules. The EJB deployment tool runs during installation of EJB 1.x or 2.x modules. The EJB deployment tool does not run during installation of EJB 3.x modules.
Map shared libraries	On the Shared library references and Shared library mapping pages, specify shared library files for your application or web modules to use. A defined shared library must exist to associate your application or module to the library file.
Map shared library relationships	On the Map shared library relationships page, specify relationship identifiers and composition unit names for shared libraries that modules in your enterprise application reference. When installing your enterprise application, the product creates a composition unit for each shared library relationship in the business-level application that you specified for Business-level application name on the Select installation options page.

Table 19. Wizard page descriptions (continued). The table describes each wizard page.

Page	Description
Provide JSP reloading options for web modules	If your application uses web modules, for Provide JSP reloading options for web modules , configure the class reloading of JavaServer Pages (JSP) files.
Map context roots for web modules	If your application uses web modules that are defined in the application XML deployment descriptor, for Map context roots for web modules , specify a context root for each web module in the application. The product does not include web modules from annotations on this page.
Initialize parameters for servlets	If your application uses web modules that support Servlet 2.5, for Initialize parameters for servlets , specify or override initial parameters that are passed to the init method of web module servlet filters. This page shows servlets from the module XML deployment descriptor. Servlet deployment information from annotations is not available on this page.
Map environment entries for EJB modules	If your application uses EJB modules, for Map environment entries for EJB modules , configure the environment entries of EJB modules such as entity, session, or message driven beans.
Map environment entries for client modules	If you are deploying one or more application client modules, for Map environment entries for client modules , configure the environment entries of client modules that are deployed as JAR files. To view the Map environment entries for client modules page, select the Deploy client modules option on the Select installation options page.
Map environment entries for web modules	If your application uses web modules that support Servlet 2.5, for Map environment entries for web modules , configure the environment entries of web modules such as servlets and JSP files.
Map environment entries for application level	If your application defines one or more environment entries, for Map environment entries for application level , configure the environment entries of applications that are deployed as EAR files.
Map resource environment entry references to resources	If your application contains resource environment references, for Map resource environment entry 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, after you click Finish a validation error is displayed.
Correct use of system identity	If your application defines Run-As Identity as <i>System Identity</i> , for Correct use of system identity , you can optionally change it to <i>Run-As role</i> and specify a user name and password for the Run As role specified. Selecting <i>System Identity</i> 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.
Correct isolation levels for all resource references	If your application has resource references that map to resources that have an Oracle database doing backend processing, for Correct isolation levels for all resource references , 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.
Map JASPI Provider	On the Map JASPI Provider page, if your application has web modules, you can specify values to override the JASPI settings from the global or domain security configuration. By default, an application inherits the JASPI settings defined in the WebSphere Application Server global or domain security configuration, and web modules inherit the application setting.

Table 19. Wizard page descriptions (continued). The table describes each wizard page.

Page	Description
Bind message destination references to administered objects	<p>If your application uses message driven beans, for Bind message destination references to administered objects, specify the JNDI name of the J2C administered object to bind the message destination reference to the message driven beans.</p> <p>If the message destination reference is from an EJB 3.0 or later module, then the JNDI name is optional and the run time provides a container default value.</p> <p>Attention: If multiple message destination references link to the same message destination, only one JNDI name is collected. When a message destination reference links to the same message destination as a message driven bean and the destination JNDI name has been collected already, the destination JNDI name for the message destination reference is not collected.</p>
Provide JNDI names for JCA objects	<p>If your application contains an embedded .rar file, for Provide JNDI names for JCA objects, specify the name and JNDI name of each JCA connection factory, administered object and activation specification.</p>
Bind J2C activationspecs to destination JNDI names	<p>If your application contains an embedded .rar file, its activationSpec property has the value <code>Destination</code>, and its introspected type is <code>javax.jms.Destination</code>, for Bind J2C activationspecs to destination JNDI names, specify the <code>jndiName</code> value for each activation bound to it.</p>
Select current backend ID	<p>If your application has EJB modules for which deployment code has been generated for multiple backend databases using an assembly tool, for Select current backend ID, specify the backend ID representing the backend database to be used when the EJB module runs.</p> <p>For information about backend databases, see topics on the EJB deployment tool.</p> <p>This step is not shown if the Deploy enterprise beans setting is enabled on the Select installation options page and if a database type other than <code>None</code> is specified on the Provide options to perform the EJB Deploy page.</p>
Metadata for modules	<p>If your application has EJB 3.x or Web 2.5 modules, you can lock deployment descriptors for one or more of the EJB 3.x or Web 2.5 modules. If you set the <code>metadata-complete</code> attribute to <code>true</code> and lock deployment descriptors, the product writes the complete module deployment descriptor, including deployment information from annotations, to XML format.</p>
Provide options to perform the web services deployment	<p>If the Deploy web services setting is enabled on the Select installation options page and your application uses web services, then you can specify <code>wsdeploy</code> command options on the Provide options to perform the web services deployment page. For information about this page, refer to descriptions of the <code>wsdeploy -cp</code> and <code>-jardir</code> options.</p>
Configure remote request dispatch properties	<p>If you are using a remote request dispatcher, you need to first configure it for use. Remote Request Dispatcher (RRD) is a pluggable extension to the web container that allows application frameworks, servlets, and JavaServer Pages to include content from outside of the currently running resource Java virtual machine (JVM) as part of the response sent to the client. Refer to Remote dispatcher property settings for information about where to configure these properties in the administration console.</p>
Display module build ID	<p>If the <code>MANIFEST.MF</code> file of a module in an enterprise application specifies a build identifier, this page shows the build identifier of the module.</p>

5. On the Summary page, verify the cell, node, and server onto which the application modules will install:
 - a. Beside **Cell/Node/Server**, click **Click here**.
 - b. Verify the settings.
 - c. Return to the Summary page.

- d. Click **Finish**.

Results

Several messages are displayed, indicating whether your application file is installing successfully.

If **Validate input off/warn/fail** on the **Select installation options** page is set to **warn**, the default, several validation warnings might be displayed. If the setting is **fail**, the validation warnings might cause errors.

If you receive an OutOfMemory error 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 error, 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 error, check the options you specified on the Java virtual machine page of the application server running the administrative console. You might increase the maximum heap size. Then, try installing the application file again.

If you are using the application installation wizard in an Internet Explorer browser, the application installation might fail intermittently. The Internet Explorer browser does not send all encrypted data expected by the server. To install the application, use the application installation wizard in a Firefox 1.5 or later browser or use the wsadmin tool.

On z/OS operating systems, application files larger than 10 megabytes might not install. When the application server `protocol_http_large_data_inbound_buffer` custom property is set to zero (0), the default, you cannot install application files larger than 10 megabytes. Set `protocol_http_large_data_inbound_buffer` to a value larger than the application file size. For example, if the EAR file that you want to install is 15 megabytes in size, set `protocol_http_large_data_inbound_buffer` to 20 megabytes:

1. Click **Servers > Server Types > WebSphere application servers > *server_name* > Server infrastructure > Administration > Custom properties**.
2. Set the `protocol_http_large_data_inbound_buffer` application server custom property to 20000000.
3. Stop and restart the application server.
4. Try installing the EAR file again.

What to do next

After the application file installs successfully, do the following:

1. Save the changes to your configuration.

For example, click the **Save** link in the application installation messages.

The application is registered with the administrative configuration and application files are copied to the target directory, which is `app_server_root/installedApps/cell_name` by default or the directory that you designate.

For a single-server product, application files are copied to the destination directory when the changes are saved.

For a multiple-server product, files are copied to remote nodes when the configuration on the deployment manager synchronizes with the configuration on individual nodes.

If you clicked the **Save** link in the application installation messages, the Preparing for the application installation page displays again. Click **Applications > Application Types > WebSphere enterprise applications** to exit the page and to see your application in the list of installed applications.

2. Start the application.
3. 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, edit the application configuration.

Example: Installing an EAR file using the default bindings

If application bindings were not specified for all enterprise beans or resources in an enterprise application during application development or assembly, you can select to generate default bindings. After application installation, you can modify the bindings as needed using the administrative console.

Before you begin

This topic assumes that the application can run on a web server.

About this task

This topic describes how to install a simple .ear file using the default bindings. You can follow the steps to install any application, including applications provided from the Samples information center.

Procedure

1. Click **Applications > New Application > New Enterprise Application** in the console navigation tree.
2. On the first Preparing for application install page, specify the full path name of the EAR file.
 - a. For **Path to the new application**, specify the full path name of the .ear file. For this example, the base file name is `my_app1.ear` and the file resides on a server at `C:/sample_apps`.
Optionally, select **Remote file system** and click **Browse**. On the Browse Remote Filesystems page, select the node that runs on the server which holds `my_app1.ear` and the EAR file name. Click **C:/, sample_apps, my_app1.ear**, and then **OK**.
 - b. Click **Next**.
3. On the second Preparing for application install page, choose to generate default bindings.
 - a. Expand **Choose to generate default bindings and mappings**.
 - b. Select **Generate default bindings**.
Using the default bindings causes any incomplete bindings in the application to be filled in with default values. The product does not change existing bindings. By choosing this option, you can skip many of the steps of the application installation wizard and go directly to the Summary step.
 - c. Click **Next**.
4. If application security warnings are displayed, read the warnings and click **Continue**.
5. On the Install New Application page, click the step number for **Map modules to servers**, and verify the cell, node, and server onto which the application files will install.
 - a. From the **Clusters and servers** list, select the server onto which the application files will install.
 - b. Select all of the application modules.
 - c. Click **Next**.
On the **Map modules to servers** page, you can map modules to other servers such as web servers. If you want a web server to serve the application, use the **Ctrl** key to select an application server or cluster and the web server together in order to have the plug-in configuration file `plugin-cfg.xml` for that web server generated based on the applications which are routed through it.
6. On the Install New Application page, click the step number beside **Summary**, the last step.
7. On the Summary page, click **Finish**.

What to do next

Examine the application installation progress messages. If the application installs successfully, save your administrative configuration. You can now see the name of your application in the list of deployed applications on the Enterprise applications page accessed by clicking **Applications > Application Types > WebSphere enterprise applications** in the console navigation tree.

If the application does not install successfully, read the messages to identify why the installation failed. Correct problems with the application as needed and try installing the application again.

If the application has a web module, try opening a browser on the application.

1. Point a web browser at the URL for the deployed application.
The URL typically has the format `http://host_name:9060/web_module_name`, where *host_name* is your valid web server and 9060 is the default port number.
2. Examine the performance of the application.

If the application does not perform as desired, edit the application configuration, then save and test it again.

Example: Installing a web services sample with the console

The product provides a web services sample application that you can install on an application server.

Before you begin

Download and extract the JaxWSServicesSample sample application. Ensure that your product installation has a Version 7.x or later application server onto which you can install the Web Services Sample.

About this task

The JaxWSServicesSamples.ear enterprise application and supporting Java archives (JAR) files are located in the `installableApps` directory within the JaxWSServicesSamples sample application.

This topic describes how to install and start the JaxWSServicesSamples.ear enterprise application using an administrative console.

Procedure

1. Click **Applications > New Application > New Enterprise Application** in the console navigation tree.
2. On the first Preparing for the application installation page, specify to install JaxWSServicesSamples.ear.
 - a. Click **Local file system** or **Remote file system** and specify the full path name of the JaxWSServicesSamples.ear file.
`../installableApps/JaxWSServicesSamples.ear`
 - b. Click **Next**.
3. On the second Preparing for the application installation page, select the fast path option.
 - a. Select **Fast Path - Prompt only when additional information is required**.
 - b. Click **Next**.
4. Click **Next** on each page until you reach the Summary page.
Do not go directly from Step 1 to the Summary page. You must click **Next** on each page that has mandatory settings to enter values for those settings. Simply click **Next** to enter the default values. You optionally can change the values to suit your environment.
5. On the Summary page, verify the cell, node, and server onto which the application modules will install, and then click **Finish**.
6. Examine the application installation progress messages.
If the application installs successfully, the message `Application JaxWSServicesSamples installed successfully` is displayed. Click **Save**. After the configuration changes are saved, you can see the name of the application in the list of deployed applications on the Enterprise applications page accessed by clicking **Applications > Application Types > WebSphere enterprise applications** in the console navigation tree.

If the application does not install successfully, read the messages to identify why the installation failed. Correct problems with the server or application and try installing the application again.

Results

The **JaxWSServicesSamples** application is in the list of deployed applications on the *Enterprise applications* page.

What to do next

After the application installs successfully, do the following:

1. Start the application.
On the Enterprise applications page, select the check boxes beside **JaxWSServicesSamples**, and then click **Start**.
2. Test the application. Point your web browser at:
`http://localhost:9080/wssamplesei/demo`
If the localhost address does not load, substitute the host name (IP address) of the computer for localhost; for example, `http://9.22.33.44:9080/wssamplesei/demo`.
If you have another WebSphere Application Server installation on your machine, the server port number is likely not 9080. See the Ports table in the administrative console to find the WC_defaulthost server port number. Click **Servers** > **Server Types** > **WebSphere application servers** > **server1** > **Ports**. The Port descriptions table lists the important ports.

Table 20. Port descriptions. Use the WC_defaulthost port in the URL to test the sample.

Port name	Description
WC_adminhost	Port used to open an unsecure administrative console in the URL <code>http://host_name:administrative_port/ibm/console</code>
WC_adminhost_secure	Port used to open a secure administrative console in the URL <code>http://host_name:administrative_port/ibm/console</code>
WC_defaulthost	Port used to test running applications in the URL <code>http://host_name:server_port/servlet_name</code>
WC_defaulthost_secure	Port used to securely test running applications in the URL <code>http://host_name:server_port/servlet_name</code>

Preparing for application installation settings

Use this page to specify an application or module to install.

To view this administrative console page, click **Applications** > **New application** > **New Enterprise Application**.

This page is the first Preparing for the application installation page. On this page, specify an application or module to install. You can install an enterprise application archive (EAR file), enterprise bean (EJB) module (JAR file), Session Initiation Protocol (SIP) module (SAR file), or web module (WAR file).

The second Preparing for the application installation page has more installation options, such as to generate default bindings for incomplete existing bindings in your application or module.

Path to the new application

Specifies the fully qualified path to the enterprise application file.

The file can be an .ear, .jar, .sar, or .war file.

During application installation, the product typically uploads application files from a client workstation running the browser to the server running the administrative console, and then deploys the application files on the server. In such cases, use the web browser running the administrative console to select EAR, WAR, SAR, or JAR modules to upload to the server.

Use **Local file system** when the browser and application files are on the same computer.

Use **Remote file system** in the following situations:

- The application file resides on any node in the current cell context. Only .ear, .jar, .sar, or .war files are shown during the browsing.
- The application file resides on the file system of any of the nodes in a cell.
- The application file already resides on the computer running the application server. For example, the field value might be `profile_root/installableApps/test.ear`.

After the product transfers the application file, the **Remote file system** value shows the path of the temporary location on the server.

Preparing for application installation binding settings

Use this page to select whether to view all installation options and to change the existing bindings for you application or module during installation. You can choose to generate default bindings for any incomplete bindings in the application or module or to assign specific bindings during installation.

This page is the second Preparing for the application installation page.

To view this administrative console page, click **Applications > New application > New Enterprise Application**, specify the path for the application or module to install, and then click **Next**.

The console page might not display all of the binding options listed in this topic. The contents of the application or module that you are installing determines which options are displayed on the console page. Also, the **Specify bindings to use** option displays only when updating an installed application.

How do you want to install the application?

Specifies whether to show only installation options that require you to supply information or to show all installation options.

Table 21. Installation option descriptions. You can select a Fast Path or select to see all installation options and parameters.

Option	Description
Fast Path - Prompt only when additional information is required	Displays only those options that require your attention, based on the contents of your application or module. Use the fast path to install your application more easily because you do not need to examine all available installation options.
Detailed - Show all installation options and parameters	Displays all available installation options.

Specify bindings to use

Specifies whether to merge bindings when you update applications or to use new or existing bindings.

This setting is shown only when you update an installed application, and not when you install a new application.

Table 22. Binding option descriptions. You can use merged, new, or existing bindings.

Option	Description
Merge new and existing bindings	The binding information from the updated application or modules is preferred over the corresponding binding information from the installed version. If any element of the binding is missing in the updated version, the corresponding element from the installed version is used. If both the installed and the updated application or module does not have a binding value, the default value is used. The product assigns a default value only if you select the Generate default bindings option.
Use new bindings	The binding information in the updated application or module is used. The binding information from the updated version of the application or module is preferred over the corresponding binding information in the installed version. The binding information from the installed version of the application or module is ignored.
Use existing bindings	The binding information from the installed version of the application or module is preferred over the corresponding binding information from the updated version. If any element of the binding information does not exist in the installed version, the element from the updated version is used. That is, bindings from the updated version of the application or module are ignored if a binding exists in the installed version. Otherwise, the new bindings are honored and not ignored.

Generate default bindings

Specifies whether to generate default bindings and mappings. To view this setting, expand **Choose to generate default bindings and mappings**. If you select **Generate default bindings**, then the product completes any incomplete bindings in the application with default values. The product does not change existing bindings.

After you select **Generate default bindings**, you can advance directly to the Summary step and install the application if none of the steps have a red asterisk (*). A red asterisk denotes that the step has incomplete data and requires a valid value. On the Summary page, verify the cell, node, and server on which the application is installed.

transition: You do not need to specify Java Naming and Directory Interface (JNDI) values for EJB bean, local home, remote home, or business interfaces of EJB 3.0 or later modules. The product assigns container default values during run time. Similarly, for any EJB reference within an EJB 3.0, EJB 3.1, Web 2.4, or Web 2.5 module, you do not need to specify JNDI values because the product resolves the targets automatically during run time. Even when you select **Generate default bindings**, the product does not generate default values for those JNDI values but it does generate default values for other bindings such as virtual host.

If you select **Generate default bindings**, the product generates bindings as follows:

- Enterprise bean (EJB) JNDI names are generated in 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 or in its corresponding annotation for EJB 3.0 or later modules. The product does not generate default values for enterprise beans in an EJB 3.0 or later module because the run time provides container default values.
- EJB references are bound if an `<ejb-link>` is found. 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. The product does not generate default values for EJB reference in an EJB 3.0, EJB 3.1, Web 2.4, or Web 2.5 module because the run time provides container default values or automatically resolves the target references.
- Resource reference bindings are derived from the `<res-ref-name>` tag or its corresponding annotation for Java Platform, Enterprise Edition (Java EE) 5 or 6 modules. 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 and EJB 2.1 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 and EJB 2.1 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 and password options. This action results in default data source settings for each JAR file. No bean-level data source bindings are generated.
- For EJB 2.0 or later message-driven beans deployed as Java EE Connector Architecture (JCA) 1.5-compliant resources, the JNDI names corresponding to activationSpec instances are generated in the form `eis/ MDB_ejb-name`. Message destination references are bound if a `<message-destination-link>` is found, then the JNDI name is set to `ejs/message-destination-linkName`. Otherwise, the JNDI name is set to `eis/message-destination-refName`.
- For EJB 2.0 or later message-driven beans deployed against listener ports, the listener ports are derived from the message-driven bean `<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, if you experience errors, complete the following actions:

- Control the global JNDI names of one or more EJB files.
- Control data source bindings for container-managed persistence (CMP) beans. That is, you have multiple data sources and need more than one global data source.
- Map resource references to global resource JNDI names that are different from the `java:comp/env` name.

In such cases, you can change the behavior with an XML document, which is a custom strategy. Use the **Specific bindings file** setting to specify a custom strategy and see the setting description in this help file for examples.

Override existing bindings

Specifies whether generated bindings are to replace existing bindings.

The default is to not override existing bindings. Select **Override existing bindings** to have generated bindings replace existing bindings.

Override existing bindings is similar to the `-defaultbinding.force` scripting option.

Specific bindings file

Specifies a bindings file that overrides the default binding.

Specific bindings file is similar to the `-defaultbinding.strategy.file` scripting option.

Change the behavior of the default binding with an XML document, which is a custom strategy. Custom strategies extend the default strategy so you only need to customize those areas where the default strategy is insufficient. Thus, 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.

Use the following examples to override various aspects of the default bindings generator:

Controlling an EJB JNDI name

```
<?xml version="1.0"?>
<!DOCTYPE df1tbindings SYSTEM "df1tbindings.dtd">
<df1tbindings>
  <module-bindings>
    <ejb-jar-binding>
      <jar-name>helloEjb.jar</jar-name>
      <ejb-bindings>
        <ejb-binding>
          <ejb-name>HelloEjb</ejb-name>
          <jndi-name>com/acme/ejb/HelloHome</jndi-name>
        </ejb-binding>
      </ejb-bindings>
    </ejb-jar-binding>
  </module-bindings>
</df1tbindings>
```

```

    </ejb-bindings>
  </ejb-jar-binding>
</module-bindings>
</df1tbndngs>

```

Remember: Ensure that the setting for <ejb-name> matches the ejb-name entry in the EJB JAR deployment descriptor. Here the setting is <ejb-name>HelloEjb</ejb-name>.

Setting the connection factory binding for an EJB JAR file

```

<!DOCTYPE df1tbndngs SYSTEM "df1tbndngs.dtd">
<df1tbndngs>
  <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>
</df1tbndngs>

```

Setting the connection factory binding for an EJB file

```

<?xml version="1.0">
<!DOCTYPE df1tbndngs SYSTEM "df1tbndngs.dtd">
<df1tbndngs>
  <module-bindings>
    <ejb-jar-binding>
      <jar-name>yourEjb20.jar</jar-name>
      <ejb-bindings>
        <ejb-binding>
          <ejb-name>YourCmp20</ejb-name>
          <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>
</df1tbndngs>

```

Restriction: Ensure that the setting for <ejb-name> matches the ejb-name tag in the deployment descriptor. Here the setting is <ejb-name>YourCmp20</ejb-name>.

Setting the message destination reference JNDI for a specific enterprise bean

This example shows an XML extract in a custom strategy file for setting message-destination-refs for a specific enterprise bean.

```

<?xml version="1.0">
<!DOCTYPE df1tbndngs SYSTEM "df1tbndngs.dtd">
<df1tbndngs>
  <module-bindings>
    <ejb-jar-binding>
      <jar-name>yourEjb21.jar</jar-name>
      <ejb-bindings>
        <ejb-binding>
          <ejb-name>YourSession21</ejb-name>
          <message-destination-ref-bindings>
            <message-destination-ref-binding>
              <message-destination-ref-name>jdbc/MyDataSrc</message-destination-ref-name>
              <jndi-name>eis/somA0</jndi-name>
            </message-destination-ref-binding>
          </message-destination-ref-bindings>
        </ejb-binding>
      </ejb-bindings>
    </ejb-jar-binding>
  </module-bindings>
</df1tbndngs>

```

Restriction: Ensure that the setting for <ejb-name> matches the ejb-name tag in the deployment descriptor. Here the setting is <ejb-name>YourSession21</ejb-name>. Also ensure that the setting for <message-destination-ref-name> matches the message-destination-ref-name tag in the deployment descriptor. Here the setting is <message-destination-ref-name>jdbc/MyDataSrc</message-destination-ref-name>.

Overriding a resource reference binding from a WAR, EJB JAR file, or Java EE client JAR file

This example shows code for overriding a resource reference binding from a WAR file. Use similar code to override a resource reference binding from an enterprise bean (EJB) JAR file or a Java EE client JAR file.

```
<?xml version="1.0"?>
<!DOCTYPE dfltbndngs SYSTEM "dfltbndngs.dtd">
<dfltbndngs>
  <module-bindings>
    <war-binding>
      <jar-name>hello.war</jar-name>
      <resource-ref-bindings>
        <resource-ref-binding>
          <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>
</dfltbndngs>
```

Restriction: Ensure that the setting for `<resource-ref-name>` matches the `resource-ref` tag in the deployment descriptor. In the previous example, the setting is `<resource-ref-name>jdbc/MyDataSrc</resource-ref-name>`.

Overriding the JNDI name for a message-driven bean deployed as a JCA 1.5-compliant resource

This example shows an XML extract in a custom strategy file for overriding the Java Message Service (JMS) `activationSpec` JNDI name for an EJB 2.0 or later message-driven bean deployed as a JCA 1.5-compliant resource.

```
<?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>
          <activation-spec-jndi-name>activationSpecJNDI</activation-spec-jndi-name>
        </ejb-binding>
      </ejb-bindings>
    </ejb-jar-binding>
  </module-bindings>
</dfltbndngs>
```

Overriding the JMS listener port name for an EJB 2.0, 2.1, or 3.0 message-driven bean

This example shows an XML extract in a custom strategy file for overriding the JMS listener port name for an EJB 2.0 or later message-driven bean deployed against a listener port.

```
<?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>
```

Overriding an EJB reference binding from an EJB JAR, WAR file, or EJB file

This example shows code for overriding an EJB reference binding from an EJB JAR file. Use similar code to override an EJB reference binding from a WAR file or an EJB file.

```
<?xml version="1.0"?>
<!DOCTYPE dfltbndngs SYSTEM "dfltbndngs.dtd">
<dfltbndngs>
  <module-bindings>
```

```

<ejb-jar-binding>
  <jar-name>YourEjbJar.jar</jar-name>
  <ejb-ref-bindings>
    <ejb-ref-binding>
      <ejb-ref-name>YourEjb</ejb-ref-name>
      <jndi-name>YourEjb/JNDI</jndi-name>
    </ejb-ref-binding>
  </ejb-ref-bindings>
</ejb-jar-binding>
</module-bindings>
</df1tbdngs>

```

Specify unique prefix for beans

Specifies a string that the product applies to the beginning of generated enterprise bean JNDI names. The prefix must be unique within the cell or node.

The default is to not specify a unique prefix for beans.

Specify unique prefix for beans is similar to the scripting option `-defaultbinding.ejbjndi.prefix`.

Default bindings for EJB 1.1 CMP beans

Specifies the default data source JNDI name and other bindings for container-managed persistence (CMP) 1.1 beans.

The default is to not use default bindings for EJB 1.1 CMP beans.

If you select **Default bindings for EJB 1.1 CMP beans**, specify the JNDI name for the default data source to be used with the CMP 1.1 beans. Also specify the user name and password for this default data source.

Default bindings for EJB 1.1 CMP beans is similar to the scripting option `-defaultbinding.datasource.jndi`.

Default connection factory bindings

Specifies the default connection factory JNDI name.

The default is to not use default connection factory bindings. Select **Default connection factory bindings** to specify bindings for connection factories.

If you select **Default connection factory bindings**, specify the JNDI name for the default connection factory to be used. Also specify whether the resource authorization is for the application or container-wide.

Default connection factory bindings is similar to the scripting option `-defaultbinding.cf.jndi`.

Use default virtual host name for web and SIP modules

Specifies the virtual host for the web module (WAR file) or Session Initiation Protocol (SIP) module (SAR file).

The default is to not use default virtual host name for web or SIP modules. If you select **Use default virtual host name for web and SIP modules**, specify a default host name.

Use default virtual host name for Web and SIP modules is similar to the scripting option `-defaultbinding.virtual.host`.

Select installation options settings

Use this page to specify options for the installation of a Java Platform, Enterprise Edition (Java EE) application onto a WebSphere Application Server deployment target. Default values for the options are used if you do not specify a value. After application installation, you can specify values for many of these options from an enterprise application settings page.

To view this administrative console page, click **Applications > New application > New Enterprise Application** and then specify values as needed for your application on the Preparing for application installation pages.

The Select installation options page is the same for the application installation and update wizards.

Precompile JavaServer Pages files

Specify whether to precompile JavaServer Pages (JSP) files as a part of installation. The default is not to precompile JSP files.

For this option, install only onto a Version 8 deployment target.

If you select **Precompile JavaServer Pages files** and try installing your application onto an earlier deployment target such as Version 7, the installation is rejected. You can deploy applications to only those deployment targets that have same version as the product. If applications are targeted to servers that have an earlier version than the product, then you cannot deploy to those targets.

Data type	Boolean
Default	false

Directory to install application

Specifies the directory to which the enterprise archive (EAR) file will be installed.

By default, the EAR file is installed in the *profile_root/installedApps/cell_name/application_name.ear* directory.

Setting options include the following:

- Do not specify a value and leave the field empty.

The default value is *\${APP_INSTALL_ROOT}/cell_name*, where the *\${APP_INSTALL_ROOT}* variable is *profile_root/installedApps*. A directory having the EAR file name of the application being installed is appended to *\${APP_INSTALL_ROOT}/cell_name*. Thus, if you do not specify a directory, the EAR file is installed in the *profile_root/installedApps/cell_name/application_name.ear* directory.

- Specify a directory.

If you specify a directory for **Directory to install application**, the application is installed in *specified_path/application_name.ear* directory. A directory having the EAR file name of the application being installed is appended to the path that you specify for **Directory to install application**. For example, if you are installing *Clock.ear* and specify *C:/myapps* on Windows computers, the application is installed in the *myapps/Clock.ear* directory. The *\${APP_INSTALL_ROOT}* variable is set to the specified path.

- Specify *\${APP_INSTALL_ROOT}/\${CELL}* for the initial installation of the application.

If you intend to export the application from one cell and later install the exported application on a different cell, specify the *\${CELL}* variable for the initial installation of the application. For example, specify *\${APP_INSTALL_ROOT}/\${CELL}* for this setting. Exporting the application creates an enhanced EAR file that has the application and its deployment configuration. The deployment configuration retains the cell name of the initial installation in the destination directory unless you specify the *\${CELL}* variable. Specifying the *\${CELL}* variable ensures that the destination directory has the current cell name, and not the original cell name.

Important: If an installation directory is not specified when an application is installed on a single-server configuration, the application is installed in *\${APP_INSTALL_ROOT}/cell_name*. When the server is made a part of a multiple-server 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

`${APP_INSTALL_ROOT}/cell_name`. However, an application that is installed after the server is added to the network configuration uses the default installation directory `${APP_INSTALL_ROOT}/network_cell_name`. To move the application to the `${APP_INSTALL_ROOT}/network_cell_name` location upon running the `addNode` operation, explicitly specify the installation directory as `${APP_INSTALL_ROOT}/${CELL}` during installation. In such a case, the application files can always be found under `${APP_INSTALL_ROOT}/current_cell_name`.

- If the application has been exported and you are installing the exported EAR file in a different cell or location, specify `${APP_INSTALL_ROOT}/cell_name/application_name.ear` if you did not specify `${APP_INSTALL_ROOT}/${CELL}` for the initial installation.

The exported EAR file is an enhanced EAR file that has the application and its deployment configuration. The deployment configuration retains the value for **Directory to install application** that was used for the previous installation of the application. Unless you specify a different value for **Directory to install application** for this installation, the enhanced EAR file will be installed to the same directory as for the previous installation.

If you did not specify the `{CELL}` variable during the initial installation, the deployment configuration uses the cell name of the initial installation in the destination directory. If you are installing on a different cell, specify `${APP_INSTALL_ROOT}/cell_name/application_name.ear`, where *cell_name* is the name of the cell to which you want to install the enhanced EAR file. If you do not designate the current cell name, *cell_name* will be the original cell name even though you are installing the enhanced EAR file on a cell that has a different name.

- Specify an absolute path or a use pathmap variable.

You can specify an absolute path or use a pathmap variable such as `{MY_APPS}`. You can use a pathmap variable in any installation.

A pathmap variable 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, accessed by clicking **Environment > WebSphere variables**.

This **Directory to install application** field is the same as the **Location (full path)** setting on an Application binaries page.

Data type	String
Units	Full path name

Distribute application

Specifies whether the product expands application binaries in the installation location during installation and deletes application binaries during uninstallation. The default is to enable application distribution. Application binaries for installed applications are expanded to the directory specified.

On single-server products, the binaries are deleted when you uninstall and save changes to the configuration.

On multiple-server products, the binaries are deleted when you uninstall and save changes to the configuration and synchronize changes.

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.

Note: If you disable this option and you do not copy and expand the application binaries to the nodes, a later saving of the configuration or manual synchronization does not move the application binaries to the nodes for you.

This **Distribute application** field is the same as the **Enable binary distribution, expansion and cleanup post uninstallation** setting on an Application binaries page.

Data type	Boolean
Default	true

Use binary configuration

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 archive (EAR) file. Select this setting for applications installed on Version 6.0 or later deployment targets only.

The default (false) is to use the binding, extensions, and deployment descriptors located in `deployment.xml`. To use the binding, extensions, and deployment descriptors located in the EAR file, enable this setting (true).

This **Use binary configuration** field is the same as the **Use configuration information in binary** setting on an Application binaries page.

Data type	Boolean
Default	false

Deploy enterprise beans

Specifies whether the EJBDeploy tool runs during application installation.

The tool generates code needed to run Enterprise JavaBeans (EJB) files. You must enable this setting in the following situations:

- The EAR file was assembled using an assembly tool such as Rational Application Developer and the EJBDeploy tool was not run during assembly.
- The EAR file was not assembled using an assembly tool such as Rational Application Developer.
- The EAR file was assembled using versions of the Application Assembly Tool (AAT) previous to Version 5.0.

If an EJB module is packaged in a web archive (WAR), you do not need to enable this setting.

The EJB deployment tool runs during installation of EJB 1.x or 2.x modules. The EJB deployment tool does not run during installation of EJB 3.x modules.

For this option, install only onto a Version 8 deployment target.

If you select **Deploy enterprise beans** and try installing your application onto an earlier deployment target such as Version 7, the installation is rejected. You can deploy applications to only those targets that have same WebSphere version as the product. If applications are targeted to servers that have an earlier version than the product, then you cannot deploy to those targets.

Also, if you select **Deploy enterprise beans** and specify a database type on the **Provide options to perform the EJB Deploy** page, previously defined backend IDs for all of the EJB modules are overwritten by the chosen database type. To enable backend IDs for individual EJB modules, set the database type to "" (null) on the **Provide options to perform the EJB Deploy** page.

Enabling this setting might cause the installation program to run for several minutes.

Data type	Boolean
Default	true (false for EJB 3.0 modules)

Application name

Specifies a logical name for the application. An application name must be unique within a cell and cannot contain an unsupported character.

An application name cannot begin with a period (.), cannot contain leading or trailing spaces, and cannot contain any of the following characters:

Table 23. Characters that you cannot use in a name. The product does not support these characters in a name.

Unsupported characters		
/ forward slash	\$ dollar sign	' single quote mark
\ backslash	= equal sign	" double quote mark
* asterisk	% percent sign	vertical bar
, comma	+ plus sign	< left angle bracket
: colon	@ at sign	> right angle bracket
; semi-colon	# hash mark	& ampersand (and sign)
? question mark]]> No specific name exists for this character combination	

This **Application name** field is the same as the **Name** setting on an Enterprise application settings page.

Data type String

Create MBeans for resources

Specifies whether to create MBeans for resources such as servlets or JSP files within an application when the application starts. The default is to create MBeans.

This field is the same as the **Create MBeans for resources** setting on a Startup behavior page.

Data type Boolean

Default true

Override class reloading settings for web and EJB modules

Specifies whether the product run time detects changes to application classes when the application is running. If this setting is enabled and if application classes are changed, then the application is stopped and restarted to reload updated classes.

The default is not to enable class reloading.

This field is the same as the **Override class reloading settings for web and EJB modules** setting on a Class loading and update detection page.

Data type Boolean

Default false

Reload interval in seconds

Specifies the number of seconds to scan the application's file system for updated files. The default is the value of the reloading interval attribute in the IBM extension (META-INF/ibm-application-ext.xmi) file of the EAR file.

The reloading interval attribute takes effect only if class reloading is enabled.

To enable reloading, specify a value greater than zero (for example, 1 to 2147483647). To disable reloading, specify zero (0). The range is from 0 to 2147483647.

This **Reload interval in seconds** field is the same as the **Polling interval for updated files** setting on a Class loading and update detection page.

Data type	Integer
Units	Seconds
Default	3

Note: For IBM extension and binding files, the .xmi or .xml file name extension is different depending on whether you are using a pre-Java EE 5 application or module or a Java EE 5 or later application or module. An IBM extension or binding file is named `ibm-*-ext.xmi` or `ibm-*-bnd.xmi` where `*` is the type of extension or binding file such as `app`, `application`, `ejb-jar`, or `web`. The following conditions apply:

- For an application or module that uses a Java EE version prior to version 5, the file extension must be `.xmi`.
- For an application or module that uses Java EE 5 or later, the file extension must be `.xml`. If `.xmi` files are included with the application or module, the product ignores the `.xmi` files.

However, a Java EE 5 or later module can exist within an application that includes pre-Java EE 5 files and uses the `.xmi` file name extension.

The `ibm-webservices-ext.xmi`, `ibm-webservices-bnd.xmi`, `ibm-webservicesclient-bnd.xmi`, `ibm-webservicesclient-ext.xmi`, and `ibm-portlet-ext.xmi` files continue to use the `.xmi` file extensions.

Deploy web services

Specifies 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 **Deploy** menu choice of an assembly tool or from a command line.

For this option, install only onto a Version 6.1 or later deployment target.

If you select **Deploy web services** and try installing your application onto an earlier deployment target such as Version 6.0, the installation is rejected. You can deploy applications to only those targets that have same version as the product. If applications are targeted to servers that have an earlier version than the product, then you cannot deploy to those targets.

Data type	Boolean
Default	false

Validate input off/warn/fail

Specifies whether the product examines the application references specified during application installation or updating and, if validation is enabled, warns you of incorrect references or fails the operation.

An application typically refers to resources using data sources for container managed persistence (CMP) beans or using resource references or resource environment references defined in deployment descriptors. The validation checks whether the resource referred to by the application is defined in the scope of the deployment target of that application.

Select **off** for no resource validation, **warn** for warning messages about incorrect resource references, or **fail** to stop operations that fail as a result of incorrect resource references.

This **Validate input off/warn/fail** field is the same as the **Application reference validation** setting on an Enterprise application settings page.

Data type String
Default warn

Process embedded configuration

Specifies whether the embedded configuration should be processed. An embedded configuration consists of files such as `resource.xml` and `variables.xml`. When selected or `true`, the embedded configuration is loaded to the application scope from the `.ear` file. If the `.ear` file does not contain an embedded configuration, the default is `false`. If the `.ear` file contains an embedded configuration, the default is `true`.

This setting affects installation of enhanced EAR files. An enhanced EAR file results when you export an installed application.

When `false`, an enhanced EAR file is installed like any other application and the product ignores its embedded configuration.

If you exported the application from a cell other than the current cell and did not specify the `$(CELL)` variable for **Directory to install application** when first installing the application, deselect this setting (`false`) to expand the enhanced EAR file in the `profile_root/installedApps/current_cell_name` directory. Otherwise, if this setting is selected (`true`), the enhanced EAR file is expanded in the `profile_root/installedApps/original_cell_name` directory, where `original_cell_name` is the cell on which the application was first installed. If you specified the `$(CELL)` variable for **Directory to install application** when you first installed the application, installation expands the enhanced EAR file in the `profile_root/installedApps/current_cell_name` directory.

Data type Boolean
Default false (deselected)

File permission

Specifies access permissions for application binaries for installed applications that are expanded to the directory specified.

The **Distribute application** option must be enabled to specify file permissions.

You can specify file permissions in the text field. You can also set some of the commonly used file permissions by selecting them from the multiple-selection list. List selections overwrite file permissions set in the text field.

You can set one or more of the following file permission strings in the list. Selecting multiple options combines the file permission strings.

Table 24. File permission string sets for list options. Select a list option or specify a file permission string in the text field.

Multiple-selection list option	File permission string set
Allow all files to be read but not written to	.*=755
Allow executables to execute	.*\.dll=755#.*\.so=755#.*\.a=755#.*\.sl=755
Allow HTML and image files to be read by everyone	.*\.htm=755#.*\.html=755#.*\.gif=755#.*\.jpg=755

Instead of using the multiple-selection list to specify file permissions, you can specify a file permission string in the text field. File permissions use a string that has the following format:

`file_name_pattern=permission#file_name_pattern=permission`

where `file_name_pattern` is a regular expression file name filter (for example, `.*\\.jsp` for all JSP files), `permission` provides the file access control lists (ACLs), and `#` is the separator between multiple entries of `file_name_pattern` and `permission`. If `#` is a character in a `file_name_pattern` string, use `\\#` instead.

If multiple file name patterns and file permissions in the string match a uniform resource identifier (URI) within the application, then the product uses the most stringent applicable file permission for the file. For example, if the file permission string is `.*\\.jsp=775#a.*\\.jsp=754`, then the `abc.jsp` file has file permission 754.

best-practices: Using regular expressions for file matching pattern compares an entire string URI against the specified file permission pattern. You must provide more precise matching patterns using regular expressions as defined by Java programming API. For example, suppose the following directory and file URIs are processed during a file permission operation:

Table 25. Example URIs for file permission operations. Results are shown following this table.

1	/opt/WebSphere/profiles/AppSrv01/installedApps/MyCell/MyApp.ear/MyWarModule.war
2	/opt/WebSphere/profiles/AppSrv01/installedApps/MyCell/MyApp.ear/MyWarModule.war/MyJsp.jsp
3	/opt/WebSphere/profiles/AppSrv01/installedApps/MyCell/MyApp.ear/MyWarModule.war/META-INF/MANIFEST.MF
4	/opt/WebSphere/profiles/AppSrv01/installedApps/MyCell/MyApp.ear/MyWarModule.war/WEB-INF/classes/MyClass.class
5	/opt/WebSphere/profiles/AppSrv01/installedApps/MyCell/MyApp.ear/MyWarModule.war/mydir/MyClass2.class
6	/opt/WebSphere/profiles/AppSrv01/installedApps/MyCell/MyApp.ear/MyWarModule.war/META-INF

The file pattern matching results are:

- `MyWarModule.war` does not match any of the URIs
- `.*MyWarModule.war.*` matches all URIs
- `.*MyWarModule.war$` matches only URI 1
- `.*\\.jsp=755` matches only URI 2
- `.*META-INF.*` matches URIs 3 and 6
- `.*MyWarModule.war/.*/.*\\.class` matches URIs 4 and 5

If you specify a directory name pattern for **File permissions**, then the directory permission is set based on the value specified. Otherwise, the **File permissions** value set on the directory is the same as its parent. For example, suppose you have the following file and directory structure:

`/opt/WebSphere/profiles/AppSrv01/installedApps/MyCell/MyApp.ear/MyWarModule.war/MyJsp.jsp`

and you specify the following file pattern string:

`.*MyApp.ear$=755#.*\\.jsp=644`

The file pattern matching results are:

- Directory `MyApp.ear` is set to 755
- Directory `MyWarModule.war` is set to 755
- Directory `MyWarModule.war` is set to 755

best-practices: Regardless of the operation system, always use a forward slash (/) as a file path separator in file patterns.

Access permissions specified here are at the application level. You can also specify access permissions for application binaries in the node-level configuration. The node-level file permissions specify the maximum (most lenient) permissions that can be given to application binaries. Access permissions specified here at application level can only be the same as or more restrictive than those specified at the node level.

This setting is the same as the **File permissions** field on the Application binaries page.

Data type String

Application build identifier

Specifies an uneditable string that identifies the build version of the application.

This **Application build identifier** field is the same as the **Application build level** field on the Application binaries page.

Data type String

Allow dispatching includes to remote resources

Specifies whether an application can dispatch includes to resources across web modules that are in different Java virtual machines in a managed node environment through the standard request dispatcher mechanism.

This field is the same as the **Allow dispatching includes to remote resources** field on the Remote request dispatch properties page.

Data type Boolean
Default false

Allow servicing includes from remote resources

Specifies whether an enterprise application can service an include request from an application.

This field is the same as the **Allow servicing includes from remote resources** field on the Remote request dispatch properties page.

Data type Boolean
Default false

Business-level application name

Specifies whether the product creates a new business-level application with the enterprise application that you are installing or makes the enterprise application a composition unit of an existing business-level application.

The default is to create a new business-level application with a setting value of `WebSphere:blaname=Anyasset,blaedition=BASE`. When you select to create a new business-level application from the drop-down list, the product creates a business-level application that has the same name as your enterprise application. If a business-level application with the name of your enterprise application exists already, the product does not create a new business-level application; it adds your enterprise application as a composition unit to that existing business-level application.

If you need to use the Shared library relationship and mapping settings page to specify dependency relationships on existing shared libraries in the business-level application, select the business-level

application name from the drop-down list. No shared libraries are shown in the page if you choose to create a new business-level application and a business-level application with the default name exists already.

To add your enterprise application to an existing business-level application, select an existing business-level application from the drop-down list. The product makes your enterprise application a composition unit of the existing business-level application.

Data type	String
Default	Create a new business-level application that has the same name as the enterprise application that you are installing.
	<code>WebSphere:blaname=Anyasset,blaedition=BASE</code>

Asynchronous request dispatch type

Specifies whether web modules can dispatch requests concurrently on separate threads and, if so, whether the server or client dispatches the requests. Concurrent dispatching can improve servlet response time.

If operations are dependant on each other, do not enable asynchronous request dispatching. Select **Disabled**. Concurrent dispatching might result in errors when operations are dependant.

Select **Server side** to enable the server to dispatch requests concurrently. Select **Client side** to enable the client to dispatch requests concurrently.

Data type	String
Default	Disabled

Allow EJB reference targets to resolve automatically

Specifies whether the product assigns default JNDI values for or automatically resolves incomplete EJB reference targets.

Select this option to enable EJB reference targets to resolve automatically if the references are from EJB 2.1 or earlier modules or from Web 2.3 or earlier modules. If you enable this option, the runtime container provides a default value or automatically resolves the EJB reference for any EJB reference that does not have a binding.

If you selected **Generate default bindings** on the Preparing for application installation page, then you do not need to select this option. The product generates default values.

If you select **Allow EJB reference targets to resolve automatically**, all modules in the application must share one deployment target. If you select this option and all of the application modules do not share a common server, after you click **Finish** on the Summary page, the product displays a warning message and does not install the application. You must deselect this setting before you click **Finish** to install the application.

Data type	Boolean
Default	false

Deploy client modules

Specifies whether to deploy client modules.

Select this option (set to `true`) if the file to deploy has one or more client modules and you want to configure environment entries for the client modules. Also select this option to configure resources such as EJB references, resource references, resource environment references, or message destination

references. Selecting this option enables you to view the Map environment entries for client modules page. If you are deploying the client modules to a federated node of a deployment manager (**Federated**) or to an application server (**Server Deployed**), select this option and set **Client deployment mode** to the appropriate value for the deployment target, **Federated** or **Server Deployed**.

If you select this option, install the client modules only onto a Version 8.0 deployment target.

Data type	Boolean
Default	false

Client deployment mode

Specifies whether to deploy client modules to an isolated deployment target (**Isolated**), a federated node of a deployment manager (**Federated**), or an application server (**Server Deployed**).

The choice of client deployment mode affects how `java:` lookups are handled. All Java URL name spaces (global, application, module, and component) are local in isolated client processes. The name spaces reside on a server in federated and server deployed client processes. The server or cluster chosen as a target for a client module determines where those name spaces are created. All `java:` lookups for federated or server deployed client modules are directed to the target server or cluster. The client module does not actually run in the target server or cluster. Multiple instances of the same client module will all share the component name space in the **Federated** and **Server Deployed** modes. Choosing the **Federated** mode is simply a declaration of intent to launch the client module using Java Network Launching Protocol (JNLP), but the Java Naming and Directory Interface (JNDI) mechanics of federated and server deployed modes are the same.

Data type	String
Default	Isolated

Validate schema

Specifies whether to validate the deployment descriptors against published Java EE deployment descriptor schemas. When this option is selected, the product analyzes each deployment descriptor to determine the Java EE specification version for the deployment descriptor, selects the appropriate schema, and then checks the deployment descriptor against the Java EE deployment descriptor schema. Validation errors result in error messages.

A Java EE deployment descriptor schema is also known as a *DTD*.

If you select this option, install your application or module only onto a Version 8.0 deployment target.

Data type	Boolean
Default	false

Manage modules settings

Use this page to specify deployment targets where you want to install the modules that are contained in your application. Modules can be installed on the same deployment target or dispersed among several deployment targets.

On single-server products, a deployment target can be an application server or web server.

On multiple-server products, a deployment target can be an application server, cluster of application servers, or Web server.

To view this administrative console page, click **Applications > Application Types > WebSphere enterprise applications > *application_name* > Manage modules**. This page is similar to the Map modules to servers page on the application installation and update wizards.

On this page, each **Module** must map to one or more targets, identified under **Server**. To change a mapping:

1. In the list of mappings, select each module that you want mapped to the same target or targets.
2. If your security configuration supports multiple domains, select the domain that has the desired clusters or servers or keep the **All domains** default.
3. From the **Clusters and servers** list, select one or more targets. Select only appropriate deployment targets for a module. You cannot install modules that use WebSphere Application Server Version 8.x features on a Version 7.x or 6.x target server. Similarly, you cannot install modules that use Version 7.x features on a Version 6.x target server.

Use the Ctrl key to select multiple targets. For example, to have a web server serve your application, press the Ctrl key and then select an application server and the web server together. The product generates the plug-in configuration file, `plugin-cfg.xml`, for that web server based on the applications which are routed through it.

4. Click **Apply**.

Attention: If an application is running, changing an application setting causes the application to restart. On stand-alone servers, the application restarts after you save the change. On multiple-server products, the application restarts after you save the change and files synchronize on the node where the application is installed. To control when synchronization occurs on multiple-server products, deselect **Synchronize changes with nodes** on the Console preferences page.

If you accessed this Manage modules page from a console enterprise application page for an already installed application, you can also use this page to view and manage modules in your application.

To view the values specified for a module configuration, click the module name in the list. The displayed module settings page shows the values specified. On the settings page, you can change existing configuration values and link to additional console pages that assist you in configuring the module.

To manage a module, select the module name in the list and click a button:

Button	Resulting action
Remove	<p>Removes the selected module from the deployed application. The module is deleted from the application in the configuration repository and also from all of the nodes where the application is installed and running or expected to run.</p> <p>On multiple-server products, if the application is running on a node when the module file is deleted from the node as a result of configuration synchronization then the product stops the application, deletes the module file from the file system of the node, and restarts the application.</p>
Update	<p>Opens a wizard that helps you update modules in an application. If a module has the same URI as a module already existing in the application, the new module replaces the existing module. If the new module does not exist in the application, it is added to the deployed application.</p> <p>On multiple-server products, if the application is running on a node when the module file is updated on the node as a result of configuration synchronization then the product stops the application, updates the module file on the file system of the node, and restarts the application. If the application is running on a node when the module file is added as a result of configuration synchronization then the product starts the newly added module without stopping and restarting the running application.</p>

Button	Resulting action
Remove File	<p>Deletes a file from a module of a deployed application.</p> <p>On multiple-server products, the file is also deleted from all the nodes where the module is installed after configuration is synchronized with nodes. If the application is running on a node when the module file is deleted from the node as a result of configuration synchronization then the product stops the application, deletes the module file from the file system of the node, and restarts the application.</p>
Export File	<p>Accesses the Export a file from an application page, which you use to export a file of an enterprise application or module to a location of your choice.</p> <p>If the browser does not prompt for a location to store the file, click File > Save as and specify a location to save the file that is shown in the browser.</p>

Display clusters and servers in the following domain

Lists the domains that your security configuration supports. This setting displays only on multiple-server products with a security configuration that supports multiple domains.

From this list, select the domain that has the clusters or servers on which you want to deploy an application or module. To see all available deployment targets, select **All Domains**.

Selecting a domain causes the **Clusters and servers** list to show only the deployment targets that are configured in the domain. You cannot deploy the modules in an application across deployment targets that belong to different security domains.

Clusters and servers

Lists the names of available deployment targets. This list is the same for every application that is installed in the cell.

From this list, select only appropriate deployment targets for a module. You must install an application, enterprise bean (EJB) module, Session Initiation Protocol (SIP) archive (SAR), web module, or client module on a Version 8.x target under any of the following conditions:

- The module supports Java Platform, Enterprise Edition (Java EE) 6.
- The module calls an 8.x runtime application programming interface (API).
- The module uses an 8.x product feature.

You must install an application, EJB, SAR, or web module on a Version 8.x or 7.x target under any of the following conditions:

- The module supports Java EE 5.
- The module calls a 7.x runtime API.
- The module uses a 7.x product feature.

If a module supports J2EE 1.4, then you must install the module on a Version 6.x, 7.x or 8.x deployment target. Modules that call a 6.1.x API or use a 6.1.x feature can be installed on a 6.1.x, 7.x or 8.x deployment target. Modules that require 6.1.x feature pack functionality can be installed on a 6.1.x deployment target that has been enabled with that feature pack or on a 7.x or 8.x deployment target.

You can install an application or module developed for a Version 5.x product on any deployment target.

Module

Specifies the name of a module in the installed (or deployed) application.

URI

Specifies the location of the module relative to the root of the application (EAR file).

Module type

Specifies the type of module, for example, a web module or EJB module.

This setting is shown on the Manage modules page accessed from a console enterprise application page.

Server

Specifies the name of each deployment target to which the module currently is mapped.

To change the deployment targets for a module, select one or more targets from the **Clusters and servers** list and click **Apply**. The new mapping replaces the previous mapping.

Client module settings

Use this page to configure a deployed client module.

To view this administrative console page, click **Applications > Application Types > WebSphere enterprise applications > *application_name* > Manage modules > *client_module_name***. This page is viewable only if the selected application contains a client module and the client deployment mode is a value other than `isolated`.

URI

Specifies the location of the client module relative to the root of the application.

Alternate deployment descriptor

Specifies the alternate deployment descriptor for the module as defined in the application deployment descriptor according to the Java Platform, Enterprise Edition (Java EE) specification.

Client module property settings

Use this page to configure the deployment mode of a deployed client module.

To view this administrative console page, click **Applications > Application Types > WebSphere enterprise applications > *application_name* > Client module deployment mode**. This page is viewable only if the selected application contains a client module.

Client module deployment mode

Specifies whether to deploy client modules to an isolated deployment target (**Isolated**), a federated node of a deployment manager (**Federated**), or an application server (**Server Deployed**).

Data type	String
Default	Isolated

Provide options to compile JavaServer Pages settings

Use this page to specify options to be used by the JavaServer Pages (JSP) compiler.

This administrative console page is a step in the application installation and update wizards. To view this page, you must select **Precompile JavaServer Pages files** on the **Select installations options** page. Thus, to view this page, click **Applications > New Application > New Enterprise Application > *application_path* > Next > Detailed - Show me all installation options and parameters > Next > Next or Continue > Precompile JavaServer Pages files > Next > Step: Provide options to compile JSPs**.

You can specify the JSP compiler options on this page only when installing or updating an application that contains web modules. After the application is installed, you must edit the JSP engine configuration parameters of a web module `WEB-INF/ibm-web-ext.xmi` file to change its JSP compiler options.

Note: For IBM extension and binding files, the .xmi or .xml file name extension is different depending on whether you are using a pre-Java EE 5 application or module or a Java EE 5 or later application or module. An IBM extension or binding file is named `ibm-*-ext.xmi` or `ibm-*-bnd.xmi` where `*` is the type of extension or binding file such as `app`, `application`, `ejb-jar`, or `web`. The following conditions apply:

- For an application or module that uses a Java EE version prior to version 5, the file extension must be `.xmi`.
- For an application or module that uses Java EE 5 or later, the file extension must be `.xml`. If `.xmi` files are included with the application or module, the product ignores the `.xmi` files.

However, a Java EE 5 or later module can exist within an application that includes pre-Java EE 5 files and uses the `.xmi` file name extension.

The `ibm-webservices-ext.xmi`, `ibm-webservices-bnd.xmi`, `ibm-webservicesclient-bnd.xmi`, `ibm-webservicesclient-ext.xmi`, and `ibm-portlet-ext.xmi` files continue to use the `.xmi` file extensions.

Web module

Specifies the name of a module within the application.

URI

Specifies the location of the module relative to the root of the application (EAR file).

JSP class path

Specifies a temporary class path for the JSP compiler to use when compiling JSP files during application installation. This class path is not saved when the application installation is complete and is not used when the application is running. This class path is used only to identify resources outside of the application which are necessary for JSP compilation and which will be identified by other means (such as shared libraries) after the application is installed. In network deployment configurations, this class path is specific to the deployment manager machine.

To specify that multiple web modules use the same class path:

1. In the list of web modules, select the **Select** check box beside each web module that you want to use a particular class path.
2. Expand **Apply Multiple Mappings**.
3. Specify the desired class path.
4. Click **Apply**.

Use full package names

Specifies whether the JSP engine generates and loads JSP classes using full package names.

When full package names are used, precompiled JSP class files can be configured as servlets in the `web.xml` file, without having to use the `jsp-file` attribute. When full package names are not used, all JSP classes are generated in the same package, which has the benefit of smaller file-system paths.

When the options `useFullPackageNames` and `disableJspRuntimeCompilation` are both `true`, a single class loader is used to load all JSP classes, even if the JSP files are not configured as servlets in the `web.xml` file.

This option is the same as the `useFullPackageNames` JSP engine parameter.

JDK source level

Specifies the source level at which the Java compiler compiles JSP Java sources. Valid values are 13, 14, and 15. The default value is 13 for pre-Java EE 5 web modules, which specifies source level 1.3 and 15 for Java EE 5 and later web modules.

Disable JSP runtime compilation

Specifies whether a JSP file should never be translated or compiled at run time, even when a `.class` file does not exist.

When this option is set to `true`, the JSP engine does not translate and compile JSP files at run time; the JSP engine loads only precompiled class files. JSP source files do not need to be present in order to load class files. You can install an application without JSP source, but the application must have precompiled class files.

For a single web application class loader to load all JSP classes, this compiler option and the **Use full package names** option both must be set to `true`.

This option is the same as the `disableJspRuntimeCompilation` JSP engine parameter.

EJB JNDI names for beans

Use this page to view and modify the Java Naming and Directory Interface (JNDI) names of non-message-driven enterprise beans in your application or module.

If your application uses Enterprise JavaBeans (EJB) 2.1 and earlier modules, on the Provide JNDI names for beans panel, specify a JNDI name for each enterprise bean in every EJB 2.1 and earlier module. You must specify a JNDI name for every EJB 2.1 and earlier enterprise bean defined in the application. For example, for the EJB module `MyBean.jar`, specify `MyBean`.

The JNDI name for an EJB module can be used for both EJB 3.x modules and pre-EJB 3.0 modules. For a pre-EJB 3.0 module, you need to provide a JNDI name for the bean. For an EJB 3.x module, you have three options

- Provide no JNDI names at all
- Select the radio button to provide a JNDI name for the bean, or
- Select the radio button to provide local or remote home JNDI names.

If no JNDI name is provided, the run time provides a default value. If JNDI name for the bean is provided, you cannot provide any JNDI name for business interface in the Provide JNDI names for business interfaces panel.

To view this administrative console page, click **Applications > Application Types > WebSphere enterprise applications > application > EJB JNDI names**.

Attention: If an application is running, changing an application setting causes the application to restart. On stand-alone servers, the application restarts after you save the change. On multiple-server products, the application restarts after you save the change and files synchronize on the node where the application is installed. To control when synchronization occurs on multiple-server products, deselect **Synchronize changes with nodes** on the Console preferences page.

Module

Specifies the name of the Enterprise JavaBeans module used by your application.

Bean

Specifies the name of an enterprise bean that is contained by the module.

URI

The Uniform Resource Identifier (URI) specifies the location of the module archive relative to the root of the application EAR.

Target Resource JNDI name

Specifies the Java Naming and Directory Interface (JNDI) name of the enterprise bean.

This is a data entry field. To modify the JNDI name bound to this bean, type the new name in this field, then select **OK**.

Data type

String

Bind EJB business settings

Use this administrative console page to specify Java Naming and Directory (JNDI) name bindings for each enterprise bean with a business interface in an EJB module. Each enterprise bean with a business interface in an EJB module must be bound to a JNDI name. For any business interface that does not provide a JNDI name, or if its bean does not provide a JNDI name, a default binding name is provided. If its bean provides a JNDI name, the default JNDI name for the business interface is provided on top of its bean JNDI name by appending the package-qualified class name of the interface.

If you specify the JNDI name for a bean in the Provide JNDI names for beans page, do not specify any business interface JNDI name in this page for the same bean. If you do not specify the JNDI name for a bean in the Provide JNDI names for beans page, you can optionally specify a business interface JNDI name. If you do not specify a business interface JNDI name, the run time provides a container default.

To view this page in the administrative console, click **Applications > Application Types > WebSphere enterprise applications > *application_name* > Bind EJB business**.

Attention: If an application is running, changing an application setting causes the application to restart. On stand-alone servers, the application restarts after you save the change. On multiple-server products, the application restarts after you save the change and files synchronize on the node where the application is installed. To control when synchronization occurs on multiple-server products, deselect **Synchronize changes with nodes** on the Console preferences page.

Module

Specifies the EJB module that contains the enterprise beans that bind to the JNDI name.

Bean

Specifies the enterprise bean that binds to the JNDI name.

URI

The Uniform Resource Identifier (URI) specifies the location of the module archive relative to the root of the application EAR.

Business Interface

Specifies the enterprise bean business interface in an EJB module.

For a no-interface view, the business interface value is an empty string ("").

JNDI Name

Specifies the JNDI name associated with the enterprise bean business interface in an EJB module.

Map default data sources for modules containing 1.x entity beans

Use this page to set the default data source mapping for EJB modules that contain 1.x container-managed persistence (CMP) beans. Unless you configure individual data sources for your 1.x CMP beans, this default mapping applies to all beans within the module.

To view this administrative console page, click **Applications > Application Types > WebSphere enterprise applications > *application_name* > Map default data sources for modules containing 1.x entity beans**.

Guidelines for using this administrative console page:

- The page displays a table that depicts the EJB modules in your application that contain 1.x CMP beans.
- Each table row corresponds to a module. A row shows the JNDI name of the data source mapping target of the EJB module *only* if you bound them together during application assembly. For every data source that is displayed, you see the corresponding security configuration.
- To set your default data source mappings:
 1. Select a row. Be aware that if you check multiple rows on this page, the data source mapping target that you select in step 2 applies to all of those EJB modules.
 2. Click **Browse** to select a data source from the new page that is displayed, the Available Resources page. The Available Resources page shows all data sources that are available mapping targets for your EJB modules.
 3. Click **Apply**. The console displays the 1.x entity bean data sources page again. In the rows that you previously selected, you now see the JNDI name of the new resource mapping target.
 4. *Before* you click **OK** to save your new configuration, set the security parameters for the data source. Use the following steps.
- To specify security settings for the default data source:
 1. Select a row. Be aware that if you check multiple rows on this page, the security settings that you select later apply to all of those data sources.
 2. Type in a user name and password that comprise the authentication alias for signing on to the data source. If these entries are not listed in the application Java Platform, Enterprise Edition (Java EE) Connector (J2C) authentication data list, you must input them into the list after saving your settings on this page. Read the information center topic on managing Java EE Connector Architecture authentication data entries for more information.
 3. Click **Apply** that immediately follows the user name and password input fields.
- Repeat all of the previous steps as necessary.
- Click **OK** to save your work.

Select

Select the check boxes of the rows that you want to edit.

EJB Module

The name of the module that contains the 1.x enterprise beans.

URI

Specifies location of the module relative to the root of the application EAR file.

JNDI name

The Java Naming and Directory Interface (JNDI) name of the default data source for the EJB module.

Data type String

User name

The user name and password that comprise the authentication alias for securing the data source.

EJB references

Use this page to view and modify the Enterprise JavaBeans (EJB) references to the enterprise beans. 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 Java Naming and Directory Interface (JNDI) name) of the resource in the target operational environment.

If your application defines EJB references, for **Map EJB references to beans**, specify JNDI names for enterprise beans that represent the logical names that are specified in EJB references. Each EJB reference defined in the application must be bound to an EJB file before clicking Finish in the Summary panel.

If the EJB reference is from an EJB 3.x, Web 2.4, Web 2.5, or Client 5.0 module, the JNDI name is optional. If the **Allow EJB reference targets to resolved automatically** option is enabled, the JNDI name is optional for all modules. The runtime provides a container default or automatically resolves the EJB reference if a binding is not provided.

To view this administrative console page, click **Applications > Application Types > WebSphere enterprise applications > application_name > EJB references**.

Values are displayed for Lookup name and EJB Link if they are configured in the application. Only one of these values is allowed. If both are set, the value must be overridden by a target resource JNDI name.

Attention: If an application is running, changing an application setting causes the application to restart. On stand-alone servers, the application restarts after you save the change. On multiple-server products, the application restarts after you save the change and files synchronize on the node where the application is installed. To control when synchronization occurs on multiple-server products, deselect **Synchronize changes with nodes** on the Console preferences page.

Module

Specifies the name of the Enterprise JavaBeans module used by your application.

Bean

Specifies the name of an enterprise bean that is contained by the module.

URI

Specifies location of the module relative to the root of the application EAR file.

Resource Reference

Specifies the name of the EJB reference that is used in the enterprise bean, if applicable, and declared in the deployment descriptor of the application module.

Class

Specifies the name of a Java class associated with this enterprise bean.

Target Resource JNDI Name

Specifies the JNDI name of the enterprise bean.

This is a data entry field. To modify the JNDI name bound to this bean, type the new name in this field, then select **OK**.

Data type

String

Resource references

Use this page to designate how the resource references of application modules map to the actual resources that are configured for the application.

To view this administrative console page, click **Applications > Application Types > WebSphere enterprise applications > application_name > Resource references**.

You can also view this page during the **Map resource references to resources** step when you install an application.

- If your application uses any of the following resource types, you can set or reset their mapping configurations:
 - Default messaging JMS queues destinations
 - Default messaging JMS topic destinations
 - Data source
 - Generic JMS connection factory
 - Mail session
 - J2C connection factory
 - JMS queue connection factory for the JMS provider of WebSphere MQ
 - JMS queue destination for WebSphere MQ
 - JMS topic connection factory for WebSphere MQ
 - JMS topic destination for WebSphere MQ
 - Unified JMS connection factory for WebSphere MQ
 - URL configuration
- The page is composed of sections that correspond to each applicable resource type. Each section heading is the class name for the resource. If your application contains only one applicable resource type, you see only one section.
- Each section contains a table. Each table row depicts a resource reference within a specific module of your application.
- The rows contain the JNDI names of resource mapping targets for your references *only* if you bound them together during application assembly. You can modify those bindings on this administrative console page.
- To set your mappings:
 1. Select a row. If you want to apply the same mapping to multiple rows, complete the steps in the section, *Set multiple JNDI names*.
 2. Click **Browse** to view a new page listing of all resources that are available mapping targets for your application references.
 3. Select a resource and click **Apply**. The console displays the Resource references page again. The JNDI name of the selected resource mapping displays in the Target Resource JNDI Name field.
 4. Repeat the previous steps as necessary.
 5. If you are editing the resource references of an existing enterprise application, click **OK**. You now return to the general configuration page for your enterprise application. If you are installing the application and have completed the **Map resource references to resources** step, continue to the next step.
- **For data sources and connection factories:** Sections for these resource types contain an additional set of steps for modifying your security settings. Use the last column in the displayed table to view the authorization type for each resource configuration per application module. You can modify the corresponding authentication method only if the authorization type is container. Container-managed authorization indicates that the product performs signon to the resource rather than the enterprise bean code. The reconfiguring process differs slightly for each authentication method option:
 - When you want to assign no authentication method to a resource:
 1. Determine which resource configurations to designate with no authentication method.
 2. Select the appropriate table rows.
 3. Click **Modify Resource Authentication Method** and select **None** from the authentication method options that are displayed above the table.
 4. Click **Apply**.
 - When you want to assign the WebSphere Application Server DefaultPrincipalMapping login configuration to a resource:

1. You must apply this option to each resource individually if you want to designate different authentication data aliases. See the topic, J2EE connector security, for more information about the default mapping configuration.
 2. Select the appropriate table rows.
 3. Click **Modify Resource Authentication Method** and select **Use default method** from the list of authentication method options that are displayed above the table.
 4. Select an authentication data entry or alias from the list.
 5. Click **Apply**.
- When you want to assign a trusted context to a resource:
 1. You must have a data source that is running at least DB2 Version 9.1 for z/OS, and the data source must have trusted context enabled.
 2. You must have a data source server that is running at least DB2 Version 9.1 for z/OS, and the data source must have trusted context enabled.
 3. Select the appropriate table rows that have trusted context enabled.
 4. Click **Modify Resource Authentication Method** and select **Use trusted connections** from the authentication method options that are displayed above the table.
 5. Select an authentication alias from the list that matches an alias that is already defined in the DB2 data source. If you do not have an alias defined that is suitable, you must define a new alias.
 6. Click **Apply**.
 7. To edit the properties of the custom login configuration, click **Mapping Properties** in the table cell.
 - When you want to assign a custom Java Authentication and Authorization Service (JAAS) login configuration to a resource:
 1. See the topic, J2EE connector security, for more information about custom JAAS login configurations.
 2. Select the appropriate table row.
 3. Click **Modify Resource Authentication Method** and select **Use custom login configuration** from the authentication method options that are displayed above the table.
 4. Select an application login configuration from the list.
 5. Click **Apply**.
 6. To edit the properties of the custom login configuration, click **Mapping Properties** in the table cell.

Set multiple JNDI names

Use this option to set the same JNDI name on multiple resources with one operation.

Click **Set multiple JNDI names** to display a menu of JNDI names. If you make a selection from this list, it is applied to the **Target Resource JNDI Name** field of all the selected rows of the table.

Modify Resource Authentication Method

Use this panel to toggle the display of a panel above the table rows.

This use of this panel is described in the **For data sources and connection factories** section.

Extended Properties

Use this panel to set additional properties on the selected resource.

Select a single table row and click **Extended Properties** to set additional properties on the selected resource. For more details on using this function, see the documentation on extending DB2 data source definitions at the application level.

Select

Select the check boxes of the rows that you want to edit.

Module

The name of a module in the application.

Bean

The name of an enterprise bean that is contained by the module.

URI

Specifies location of the module relative to the root of the application EAR file.

Resource Reference

The name of a resource reference that is used in the enterprise bean, if applicable, and is declared in the deployment descriptor of the application module.

Target Resource JNDI name

The Java Naming and Directory Interface (JNDI) name of the resource that is the mapping target of the resource reference.

Data type String

Login configuration

This column applies to data sources and connection factories only and refers to the authorization type and the authentication method for securing the resource.

Virtual hosts settings

Use this page to specify virtual hosts for web modules contained in your application. Web modules can be installed on the same virtual host or dispersed among several virtual hosts.

To view this administrative console page, click **Applications > Application Types > WebSphere enterprise applications > *application_name* > Virtual hosts**. This page is the same as the **Map virtual hosts for web modules** page on the application installation and update wizards.

On this page, each web module must map to a previously defined virtual host, identified under **Virtual host**. You can see information on previously defined virtual hosts by clicking **Environment > Virtual hosts** in the administrative console. Virtual hosts enable you to associate a unique port with a module or application. The aliases of a virtual host identify the port numbers defined for that virtual host. A port number specified in a virtual host alias is used in the URL that is used to access artifacts such as servlets and JavaServer Pages (JSP) files in a web module. For example, the alias `myhost:8080` is the `host_name:port_number` portion of the URL `http://myhost:8080/servlet/snoop`.

The default virtual host setting usually is `default_host`, which provides several port numbers through its aliases:

80 An internal, insecure port used when no port number is specified
9080 An internal port
9443 An external, secure port

Unless you want to isolate your web module from other modules or resources on the same node (physical machine), `default_host` is a suitable virtual host for your web module.

In addition to `default_host`, the product provides `admin_host`, which is the virtual host for the administrative console system application. `admin_host` is on port 9060. Its secure port is 9043. Do not select `admin_host` unless the web module relates to system administration.

To change a mapping:

1. In the list of mappings, select the **Select** check box beside each web module that you want mapped to a particular virtual host.
2. From the **Virtual host** drop-down list, select the desired virtual host. If you selected more than one virtual host in step 1:
 - a. Expand **Apply Multiple Mappings**.
 - b. Select the desired virtual host from the **Virtual Host** drop-down list.
 - c. Click **Apply**.
3. Click **OK**.

Attention: If an application is running, changing an application setting causes the application to restart. On stand-alone servers, the application restarts after you save the change. On multiple-server products, the application restarts after you save the change and files synchronize on the node where the application is installed. To control when synchronization occurs on multiple-server products, deselect **Synchronize changes with nodes** on the Console preferences page.

Web module

Specifies the name of a web module in the application that you are installing or that you are viewing after installation.

Virtual host

Specifies the name of the virtual host to which the Web module is currently mapped.

Expanding the drop-down list displays a list of previously defined virtual hosts. To change a mapping, select a different virtual host from the list.

Do not specify the same virtual host for different web modules that have the same context root and are deployed on targets belonging to the same node even if the web modules are contained in different applications. Specifying the same virtual host causes a validation error.

Security role to user or group mapping

Use this page to specify the users and groups that are mapped to the security roles that are used with the enterprise application.

To view this administrative console page, click **Applications > Application types > WebSphere enterprise applications > application_name**. Under Detail Properties, click **Security role to user/group mapping**.

Table 26. User and group mapping. User and group mapping.

Button	Resulting action
Map Users	Lists the users that are mapped to the specified role within this application. If trusted realms are configured, a drop-down list of realms to search is displayed. Users from the non-default realm are displayed as user@realm
Map Groups	Lists the groups that are mapped to this specified role within this application. If trusted realms are configured, a drop-down list of realms to search is displayed. Users from the non-default realm are displayed as user@realm

Table 26. User and group mapping (continued). User and group mapping.

Button	Resulting action
Map Special Subjects	<p>This choice appears if multiple realms are being used. It enables you to map any of the following Special Subjects to a selected role:</p> <ul style="list-style-type: none"> • All authenticated in application realm: All authenticated users that are in the applications realm, which specifies whether to map all of the authenticated users to a specified role. When you map all authenticated users to a specified role, all of the valid users in the current registry who have been authenticated can access resources that are protected by this role. This selection also applies to all authenticated users regardless of the realm. • Everyone: map everyone to the selected role. When you map everyone to a role, anyone can access the resources that are protected by this role and, essentially, there is no security. • None: Do not map anyone to the selected role <p>Attention:</p> <ul style="list-style-type: none"> • If the secured realm cannot be reached, the left list is replaced with 3 text fields (that is, name, realm, and uid). You can add the user when the secured realm is not available. It is not possible to map two subjects to the same role in this release of WebSphere Application Server.

Role

Lists the specific capabilities to a user. Role privileges give users and groups permission to run as specified.

For example, you might map the user Joe to the administrator role, which enables user Joe to perform all of the tasks associated with the administrator role.

The authorization policy is only enforced when global security is enabled.

Mapped users

Lists the users that are mapped to the specified role within this application.

Special subjects

Lists which special subjects are mapped to the security role when an application uses multiple realms.

Mapped groups

Lists the groups that are mapped to this specified role within this application.

JASPI authentication enablement for applications

Use this page to enable or disable Java Authentication SPI (JASPI) authentication for an application or web module, and to specify the name of a JASPI authentication provider to be used for authenticating messages for the application or web module.

To view this administrative console page, click **Applications > Application Types > WebSphere enterprise applications**. Select an application, and under Detail Properties, select **JASPI provider**.

Select JASPI provider

Select to indicate the web modules in the application that you wish to specify or to override the default JASPI authentication settings for.

Select one of the JASPI provider names to choose a provider to use to perform authentication of web requests for the selected Web module or the application.

To specify how JASPI authentication is performed for the selected web module or the application, choose one of the following:

Do not use JASPI

Select to disable JASPI authentication for the selected web module or for the application.

Inherit JASPI provider

Select to inherit the JASPI authentication settings from default values in the cell or domain security configuration, as appropriate.

When Inherit JASPI provider is selected for a web module, JASPI authentication settings for the selected module are the settings that are specified for the application.

When Inherit JASPI provider is selected for the application, JASPI authentication settings are the settings that are specified in the appropriate cell or domain security configuration.

Provider name

When a specific provider name is selected, that provider is used to perform authentication of web requests for the selected application or web module.

If JASPI authentication is enabled, and a specific provider name is not specified, then the default provider name is used. For more information, read about configuring a new JASPI authentication provider using the administrative console.

If JASPI authentication is disabled, or if no default provider is selected, JASPI authentication is not performed. Web authentication is then performed according to another authentication mechanism as selected in the cell or domain security configuration.

User RunAs collection

Use this page to map a specified user identity and password to a RunAs role. This panel enables you to specify application-specific privileges for individual users to run specific tasks using another user identity.

To view this administrative console page, complete the following steps:

1. Click **Applications > Enterprise applications > *application_name***.
2. Under Detail properties, click **User runAs roles**.

The enterprise beans that you install contain predefined RunAs roles. RunAs roles are used by enterprise beans that need to run as a particular role for recognition while interacting with another enterprise bean.

Username

Specifies a user name for the RunAs role user.

This user already maps to the role specified in the Mapping users and groups to roles panel. You can map the user to its appropriate role by either mapping the user to that role directly or mapping a group that contains the user to that role. After you specify the user name and password for the user and select a RunAs role, click **Apply**.

Data type: String

Password

Specifies the password for the RunAs user.

Data type: String

Role

Maps specific capabilities to a user.

The authorization policy is only enforced when global security is enabled.

Ensure all unprotected 1.x methods have the correct level of protection

Use this page to verify that the unprotected Enterprise JavaBeans (EJB) Version 1.x methods have the correct level of protection before you map users to roles.

This administrative console panel is displayed during the application deployment process. To access the administrative console panel, click **Application > New application > New Enterprise Application** . The panel is displayed as *Ensure all unprotected 1.x methods have the correct level of protection* in the application deployment steps. On this administrative console panel, you can specify whether users can access specific EJB modules.

EJB module

Specifies the EJB module name.

URI

Specifies the Uniform Resource Identifier (URI) that is used to locate the Java archive (JAR) file for the EJB module.

Deny all access

Select this option to protect this EJB module by making it inaccessible to users regardless of their access permissions.

Default: Cleared

Bind listeners for message-driven beans settings

Use this page to specify bindings for message-driven beans in your application or module.

To view this administrative console page, click **Applications > Application Types > WebSphere enterprise applications > application_name > Message Driven Bean listener bindings**. This page is the same as the **Bind listeners for message-driven beans** page on the application installation and update wizards.

Each message-driven bean must be bound to a listener port name or to an activation specification Java Naming and Directory Interface (JNDI) name.

Provide a listener port name if your application uses any of the following Java Message Service (JMS) providers:

- V5 default messaging provider
- WebSphere MQ messaging provider
- Generic messaging provider

Provide an activation specification JNDI name if your application's resources are configured using the default messaging provider or any generic J2C resource adapter that supports inbound messaging.

Not providing valid listener port names or activation specification JNDI names results in the following errors:

- If neither a listener port name or an activation specification JNDI name is specified for a message driven bean, then a validation error is displayed after you click **Finish** on the **Summary** page.

- If multiple message driven beans are linked to the same destination, specify the same destination JNDI name for each message driven bean. If you specify different destination JNDI names, a validation error is displayed and all JNDI specifications after the first one are ignored.

To apply binding changes to multiple mappings:

1. In the list of mappings, select the **Select** check box beside each EJB module that you want mapped to a particular binding.
2. Expand **Apply Multiple Mappings**.
3. Complete one of the following steps:
 - Specify a listener port name.
 - Select a target resource JNDI name for an activation specification. Optionally specify the following parameters:

Destination JNDI name
For resource adapters that support JMS, specify `javax.jms.Destinations` so the resource adapter can service messages from the JMS destination. A destination JNDI name set as part of application deployment take precedence over properties set on an activation specification administrative object.

ActivationSpec authentication alias
Specify an authentication alias that is used to access the user name and password that are set on the configured J2C activation specification. Authentication alias properties set as part of application deployment take precedence over properties set on an activation specification administrative object.
4. Click **Apply**.
5. Click **OK** or **Next**.

Module

Specifies the name of the module that contains the enterprise bean.

Bean

Specifies name of an enterprise bean in the application.

URI

Specifies the location of the module relative to the root of the application EAR file.

Messaging Type

Specifies the type of message-driven bean.

Listener Bindings

Specifies a listener port name or an activation specification JNDI name for the message-driven bean. When a message-driven enterprise bean is bound to an activation specification JNDI name you can also specify the destination JNDI name and the authentication alias.

Bindings specify JNDI names for the referenceable and referenced artifacts in an application. An example JNDI name for a listener port to be used by a Store application might be `StoreMdbListener`. The binding definition is stored in IBM bindings files such as `ibm-ejb-jar-bnd.xmi`.

Note: For IBM extension and binding files, the `.xmi` or `.xml` file name extension is different depending on whether you are using a pre-Java EE 5 application or module or a Java EE 5 or later application or module. An IBM extension or binding file is named `ibm-*-ext.xmi` or `ibm-*-bnd.xmi` where `*` is the type of extension or binding file such as `app`, `application`, `ejb-jar`, or `web`. The following conditions apply:

- For an application or module that uses a Java EE version prior to version 5, the file extension must be `.xmi`.

- For an application or module that uses Java EE 5 or later, the file extension must be .xml. If .xmi files are included with the application or module, the product ignores the .xmi files.

However, a Java EE 5 or later module can exist within an application that includes pre-Java EE 5 files and uses the .xmi file name extension.

The `ibm-webservices-ext.xmi`, `ibm-webservices-bnd.xmi`, `ibm-webservicesclient-bnd.xmi`, `ibm-webservicesclient-ext.xmi`, and `ibm-portlet-ext.xmi` files continue to use the .xmi file extensions.

Map data sources for all 2.x CMP beans

Use this page to set the default data source mapping for EJB modules that contain 2.x container-managed persistence (CMP) beans. Unless you configure individual data sources for your 2.x CMP beans, this default mapping applies to all beans within the module.

To view this administrative console panel, click **Applications > Application Types > Websphere enterprise applications > *application_name* > Map data sources for all 2.x CMP beans** .

This panel displays a table that depicts the EJB modules in your application that contain 2.x CMP beans. Each table row corresponds to a module. A row shows the JNDI name of the data source mapping target of the EJB module only if you bound them together during application assembly. For every data source that is displayed, you see the corresponding security configuration.

Set Multiple JNDI Names

Specifies the JNDI name to bind to one or more modules. Select one or more modules, click **Set Multiple JNDI Names**, and select the JNDI name for the resource to which you would like to bind the module.

Set Authorization Type

Specifies the authorization type that you to use for the modules. Select one or more modules, click **Set Authorization Type**, and select the authorization type.

You can choose:

- Per application - indicates that the enterprise bean code performs signon.
- Container - indicates that the application server performs signon to the data source.

Modify Resource Authentication Method

Specifies the resource authentication method for the modules that you have configured with container-managed authorization. Select one or more modules, click **Modify Resource Authentication Method**, and select the authentication method.

You can choose between the following authentication methods:

- **None:**
 1. Determine which data source configurations to designate with no authentication method.
 2. Select the appropriate table rows.
 3. Select **None** from the list of authentication method options that precede the table.
 4. Click **Apply**.
- **Use default method (many-to-one mapping):**
 1. Determine which data source configurations to designate with the WebSphere Application Server DefaultPrincipalMapping login configuration. Apply this option to each data source individually if you want to designate different authentication data aliases. See the information center topic on J2EE Connector security for more information on the default mapping configuration.
 2. Select the appropriate table rows.

3. Select **Use default method (many-to-one mapping)** from the list of authentication method options that precede the table.
 4. Select an authentication data entry or alias from the list.
 5. Click **Apply**.
- **Use Kerberos authentication:** Specifies to use the Kerberos authentication method.
 1. Ensure that you have configured the Kerberos authentication mechanism in the application server.
 2. Select the appropriate table row.
 3. Select **Use Kerberos authentication** from the list of authentication method options that precede the table.
 4. Select an application login configuration from the list.
 5. Click **Apply**.
 6. To edit the properties of the custom login configuration, click **Mapping Properties** in the table cell.

The application server will attempt to verify that you are connecting to the correct type of database when you select this option.

- **Use trusted connections (one-to-one mapping):**
 1. Determine which data source configurations to designate with a custom Java Authentication and Authorization Service (JAAS) login configuration. See the information center topic on J2EE Connector security for more information on custom JAAS login configurations.
 2. Select the appropriate table row.
 3. Ensure that the database to which the modules will connect is configured for trusted connections.
 4. Select **Use trusted connections (one-to-one mapping)** from the list of authentication method options that precede the table.
 5. Select an application login configuration from the list.
 6. Click **Apply**.

The application server will attempt to verify that you are connecting to the correct type of database when you select this option.

- **Custom login configuration:**
 1. Determine which data source configurations to designate with a custom Java Authentication and Authorization Service (JAAS) login configuration. See the information center topic on J2EE Connector security for more information on custom JAAS login configurations.
 2. Select the appropriate table row.
 3. Select **Use custom login configuration** from the list of authentication method options that precede the table.
 4. Select an application login configuration from the list.
 5. Click **Apply**.
 6. To edit the properties of the custom login configuration, click **Mapping Properties** in the table cell.

Select

Select the check boxes of the rows you want to edit.

EJB Module

Specifies the name of the module that contains the 2.x enterprise beans.

URI

Specifies location of the module relative to the root of the application EAR file.

JNDI name

Specifies the Java Naming and Directory Interface (JNDI) name of the default data source for the EJB module.

Data type

String

Resource authorization

Specifies the authorization type and the authentication method for securing the data source.

Extended Datasource Properties

When selected, you will be directed to a panel on which you can specify extended properties that the module can use for the DB2 data source.

The application server will attempt to verify that you are connecting to the correct type of database when you select this option.

Map data sources for all 2.x CMP beans settings

Use this page to map container-managed persistence (CMP) 2.x beans of an application to data sources that are available to the application.

To view this administrative console page, click **Applications > Application Types > Websphere enterprise applications > *application_name* > Map data sources for all 2.x CMP beans.**

Each table row corresponds to a CMP bean within a specific EJB module. A row shows the JNDI name of the data source mapping target of the bean only if you bound them together during application assembly. For every data source that is displayed, you see the corresponding security configuration.

Set Multiple JNDI names

Specify the Java Naming and Directory Interface (JNDI) name for multiple EJB modules. Select one or more EJB modules from the table, and select a JNDI name from this list to configure the EJB modules with that JNDI name.

Data type

Drop-down list

Set Authorization Type

Specify the authorization type for securing the data source. Select one or more EJB modules from the table to set the authorization type.

Select either **Container** or **Application** from the displayed list. Container-managed authorization indicates that WebSphere Application Server performs signon to the data source. Application-managed authorization indicates that the enterprise bean code performs signon.

Modify Resource Authentication Method

Specify the authorization type and the authentication method for securing the data source. Select one or more EJB modules from the table to modify the resource authentication method.

You can choose between the following authentication methods:

- **None:**

1. Determine which data source configurations to designate with no authentication method.
2. Select the appropriate table rows.
3. Select **None** from the list of authentication method options that precede the table.
4. Click **Apply**.

- **Use default method (many-to-one mapping):**

1. Determine which data source configurations to designate with the WebSphere Application Server DefaultPrincipalMapping login configuration. Apply this option to each data source individually if you

want to designate different authentication data aliases. See the information center topic on J2EE Connector security for more information on the default mapping configuration.

2. Select the appropriate table rows.
 3. Select **Use default method (many-to-one mapping)** from the list of authentication method options that precede the table.
 4. Select an authentication data entry or alias from the list.
 5. Click **Apply**.
- **Use Kerberos authentication:** Specifies to use the Kerberos authentication method.
 1. Ensure that you have configured the Kerberos authentication mechanism in the application server.
 2. Select the appropriate table row.
 3. Select **Use Kerberos authentication** from the list of authentication method options that precede the table.
 4. Select an application login configuration from the list.
 5. Click **Apply**.
 6. To edit the properties of the custom login configuration, click **Mapping Properties** in the table cell.

The application server will attempt to verify that you are connecting to the correct type of database when you select this option.

- **Use trusted connections (one-to-one mapping):**
 1. Determine which data source configurations to designate with a custom Java Authentication and Authorization Service (JAAS) login configuration. See the information center topic on J2EE Connector security for more information on custom JAAS login configurations.
 2. Select the appropriate table row.
 3. Ensure that the database to which the modules will connect is configured for trusted connections.
 4. Select **Use trusted connections (one-to-one mapping)** from the list of authentication method options that precede the table.
 5. Select an application login configuration from the list.
 6. Click **Apply**.

The application server will attempt to verify that you are connecting to the correct type of database when you select this option.

- **Custom login configuration:**
 1. Determine which data source configurations to designate with a custom Java Authentication and Authorization Service (JAAS) login configuration. See the information center topic on J2EE Connector security for more information on custom JAAS login configurations.
 2. Select the appropriate table row.
 3. Select **Use custom login configuration** from the list of authentication method options that precede the table.
 4. Select an application login configuration from the list.
 5. Click **Apply**.
 6. To edit the properties of the custom login configuration, click **Mapping Properties** in the table cell.

Select

Select the check boxes of the rows that you want to edit.

EJB

The name of an enterprise bean in the application.

EJB Module

The name of the module that contains the enterprise bean.

URI

Specifies location of the module relative to the root of the application EAR file.

Target resource JNDI name

Specifies the resource to which the CMP bean is bound.

Resource authorization

Specifies the current setting for the resource authorization type.

Modify this setting with **Set authorization type**.

Ensure all unprotected 2.x methods have the correct level of protection

Use this page to verify that the unprotected Enterprise JavaBeans (EJB) Version 2.x methods have the correct level of protection before you map users to roles.

This administrative console panel is displayed during the application deployment process. To access the administrative console panel, click **Applications > New application** *application_name*. The panel is displayed as *Ensure all unprotected 2.x methods have the correct level of protection* in the application deployment steps. On this administrative console panel, you can specify whether users can access specific EJB modules.

To use this administrative console page, select the **Uncheck**, **Exclude**, or **Role** option, the check box next to the EJB module, and click **Apply**. If you select **Role** option, select the appropriate role for the EJB module before you click **Apply**.

Uncheck

Select this option if you do not want the application server to verify the access permissions for the EJB module. Everyone can access the EJB module.

Default: Selected

Exclude

Select this option to protect this EJB module by making it inaccessible to users regardless of their access permissions.

Default: Deselected

Role

Specifies the EJB level of protection based on the security role.

The roles listed in this menu are obtained from the application scope. If the selected role is not in the module, then it is added to the modules or Java archive (JAR) files.

Default: Deselected

EJB module

Specifies the name of the module.

If a module name appears in this list, then the module contains unprotected EJB methods.

URI:

Specifies the Uniform Resource Identifier (URI) that is used to locate the Java archive (JAR) file for the EJB module.

Protection type

Specifies the level of protection that is assigned to a particular module name.

After you select the **Uncheck**, **Exclude**, or **Role** option and click **Apply**, the selected protection option is displayed in this column.

Provide options to perform the EJB Deploy settings

Use this page to specify options for the enterprise bean (EJB) deployment tool. The tool generates code needed to run enterprise bean files. You can specify extra class paths, Remote Method Invocation compiler (RMIC) options, database types, and database schema names to be used while running the EJB deployment tool.

This administrative console page is a step in the application installation and update wizards. To view this page, you must select **Deploy enterprise beans** on the **Select installation options** page. Thus, to view this page, click **Applications > New Application > New Enterprise Application > *application_path* > Next > Detailed - Show all installation options and parameters > Next > Deploy enterprise beans > Next > Step: Provide options to perform the EJB Deploy**.

You can specify the EJB deployment tool options on this page when installing or updating an application that contains EJB modules. The EJB deployment tool runs during installation of EJB 1.x or 2.x modules. The EJB deployment tool does not run during installation of EJB 3.x modules.

The options that you specify set parameter values for the `ejbdeploy` command. The tool, and thus the `ejbdeploy` command, is run on the enterprise archive (EAR) file during installation after you click **Finish** on the **Summary** page of the wizard.

Class path

Specifies the class path of one or more zipped or Java archive (JAR) files on which the JAR or EAR file being installed depends.

To specify the class paths of multiple entries, the file names must be fully qualified, separated by a path separator that the target server uses, and enclosed in double quotation marks.

```
path\myJar1.jar;path\myJar2.jar;path\myJar3.jar
```

On the other supported operating systems, the path separator is a colon (:). For example:

```
path/myJar1.jar:path/myJar2.jar:path/myJar3.jar
```

Class path is the same as the `ejbdeploy` command parameter `-cp class_path`.

Data type	String
Default	null

RMIC

Specifies whether the EJB deployment tool passes RMIC options to the Remote Method Invocation compiler. Refer to RMI Tools documentation for information on the options.

Separate options by a space and enclose them in double quotation marks. For example:

```
"-nowarn -verbose"
```

The **RMIC** setting is the same as the `ejbdeploy` command parameter `-rmic "options"`.

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

Default null

Database type

Specifies the name of the database vendor, which is used to determine database column types, mapping information, `Table.sql`, and other information. Select a database type or the empty choice from the drop-down list. The list contains the names of valid database vendors. Selecting the empty choice sets the database type to "" (null).

If you specify a database type, previously defined backend IDs for all of the EJB modules are overwritten by the chosen database type. To enable backend IDs for individual EJB modules, select the empty choice to set the database type to null.

Note: The backend IDs SQL92 (1992 SQL Standard) and SQL99 (1999 SQL Standard) are deprecated. Although the SQL92 and SQL99 backend IDs are available in the list on the Provide options to perform the EJB Deploy page, they are deprecated.

Database type is the same as the `ejbdeploy` command parameter `-dbvendor name`.

Data type String
Default DB2UDB_V82

Database schema

Specifies the name of the schema that you want to create.

The EJB deployment tool saves database information in the schema document in the JAR or EAR file, which means that the options do not need to be specified again. It also means that when a JAR or EAR is generated, the correct database must be defined at that point because it cannot be changed later.

If the name of the schema contains any spaces, the entire name must be enclosed in double quotes. For example:

"my schema"

Database schema is the same as the `ejbdeploy` command parameter `-dbschema "name"`.

Data type String
Default null

Database access type

Specifies the database access type for a DB2 database that supports Structured Query Language for Java (SQLJ). Use SQLJ to develop data access applications that connect to DB2 databases. SQLJ is a set of programming extensions that support use of the Java programming language to embed statements that provide SQL (Structured Query Language) database requests.

To view this setting, you must select a DB2 backend database that supports SQLJ from the **Database type** drop-down list.

Available database access types include JDBC and SQLJ.

Data type String
Default JDBC

SQLJ class path

Specifies the class path of the DB2 SQLJ tool `sqlj.zip` file. The product uses this class path to run the DB2 SQLJ tool during application installation and generate SQLJ profiles (`.ser` files).

To view this setting, you must select a DB2 backend database that supports SQLJ from the **Database type** drop-down list.

Specify the drive and directory where the `sqlj.zip` file resides. For example:

On all other operating systems, specify `/SQLJ/sqlj.zip`.

When you reinstall an application EAR file, the product deletes any existing SQLJ profiles and creates new profiles.

If you do not specify a class path, the product displays a warning about the missing class path. After you specify a valid class path, you can continue using the wizard for the application installation.

You can customize or add bindings to the generated SQLJ profile after the product installs the application. Use the administrative console SQLJ profiles and pureQuery bind files page accessed by clicking **Applications > Application Types > WebSphere enterprise applications > *application_name* > SQLJ profiles and pureQuery bind files**.

Data type	String
Default	null

JDK compliance level

Specifies the Java developer kit compiler compliance level as *1.4*, *5.0*, or *6.0* when you include application source files for compilation.

The default is to use whatever developer kit version the `ejbdeploy` command is using. If your application is using new functionality defined in Version 5.0 or 6.0 or you are including source files (which is not recommended), then you must specify the Version 5.0 or 6.0 level.

JDK compliance level is the same as the `ejbdeploy` command parameter `-complianceLevel "1.4" | "5.0" | "6.0"`.

Data type	String
Default	null (empty string)

Shared library reference and mapping settings

Use the Shared library references and Shared library mapping pages to associate defined shared libraries with an application or web module. A shared library is an external Java archive (JAR) file that is used by one or more applications. Using shared libraries enables multiple applications deployed on a server to use a single library, rather than use multiple copies of the same library. After you associate shared libraries with an application or module, the application or module class loader loads classes represented by the shared libraries and makes those classes available to the application or module.

To view the Shared library references console page, click **Applications > Application Types > WebSphere enterprise applications > *application_name* > Shared library references**. To view the Shared library mapping page, click **Reference shared libraries** on the Shared library references page. These pages are the same as the Map shared libraries and Map shared libraries to an entire application or module pages in the application installation and update wizards.

On the Shared library references page, the first element listed is the application. The other elements are modules in the application.

To associate shared libraries with your application or module:

1. Select an application or module.
2. Click **Reference shared libraries**.
3. On the Shared library mapping page, select one or more shared libraries that the application or modules uses in the **Available** list, click >> to add them to the **Selected** list, and click **OK**.

A defined shared library for a file that your application or module uses must exist to associate your application or module to the library.

If no shared libraries are defined and the application is installed already, on the Shared library mapping page, click **New** and define a shared library.

You can otherwise define a shared library as follows:

1. Click **Environment > Shared libraries**.
2. Specify whether the shared library is visible at the cell, node or server level.
3. Click **New**.
4. On the settings page for the new shared library, specify a name and one or more class paths. If the libraries are platform-specific files such as .dll, .so, or *SRVPGM objects, also specify a native library path. Then, click **Apply**.
5. Save the administrative configuration.

Application

Specifies the name of the application that you are installing or that you selected on the Enterprise applications page.

Module

Specifies the name of the module associated with the shared libraries.

URI

Specifies the location of the module relative to the root of the application EAR file.

Shared libraries

Specifies the name of the shared library files associated with the application or module.

Shared library relationship and mapping settings

Use the Shared library relationship and Shared library relationship mapping pages to specify relationship identifiers and composition unit names for shared libraries that modules in your enterprise application reference. When installing your enterprise application, the product creates a composition unit for each shared library relationship in the business-level application that you specified on the Select installation options page of the application installation wizard.

To view this console page in a wizard, click **Applications > Install new application > New Enterprise Application > *application_path* > Next > Detailed - Show all installation options and parameters > Next > *application_name* > Step: Map shared library relationships**.

After installation, click **Applications > Application Types > WebSphere enterprise applications > Shared library relationships**.

To map library files used in a business-level application to an application or web module, use the Shared library relationship mapping page:

1. Click **Reference shared libraries**.

2. Note the application or module in **Map libraries to the application or module listed**. You are associating library files with that application or module.
3. From the **Available** list, select one or more libraries that the application or module uses.
4. Click >> to add them to the **Selected** list.
5. To remove an association, select one or more libraries in the **Selected** list and click <<.
6. Click **OK**.

Module

Specifies the name of the module associated with the shared libraries.

URI

Specifies the location of the module relative to the root of the application EAR file.

Relationship identifiers

Specifies an identifier for a module shared library relationship. The product assigns an identifier to the composition unit that it creates for the shared library relationship in the business-level application.

Composition unit names

Specifies a composition unit name for the shared library relationship. The product uses this value to name the composition unit that it creates for the shared library relationship in the business-level application that you specified on the Select installation options page of this wizard.

This setting is only in the application installation and update wizards.

Match target

Specifies whether the product maps the composition unit for the shared library relationship to the same deployment target as the business-level application.

Note: If you later change the deployment target of the business-level application or its modules, you must manually update the shared library target to match the target of the application and modules. The targets of shared library composition units are not automatically updated. Not updating the target of the shared library composition unit might cause `java.lang.ClassNotFoundException` errors and prevent the application or its modules from starting. To prevent these error conditions, also ensure that shared libraries upon which other modules or applications depend have a lower starting weight than dependent applications and modules.

JSP and JSF option settings

Use this panel to configure the class reloading of web modules such as JavaServer Pages (JSP) files and to select a JSF implementation to use with this application.

To view this administrative console panel, click **Applications > Application Types > WebSphere enterprise applications > *application_name* > JSP and JSF options**. This panel is the same as the **Provide JSP reloading options for web modules** panel on the application installation and update wizards.

The following note applies to the files with a .xmi extension in this topic:

Note: For IBM extension and binding files, the .xmi or .xml file name extension is different depending on whether you are using a pre-Java EE 5 application or module or a Java EE 5 or later application or module. An IBM extension or binding file is named `ibm-*-ext.xmi` or `ibm-*-bnd.xmi` where * is the type of extension or binding file such as `app`, `application`, `ejb-jar`, or `web`. The following conditions apply:

- For an application or module that uses a Java EE version prior to version 5, the file extension must be .xmi.

- For an application or module that uses Java EE 5 or later, the file extension must be .xml. If .xmi files are included with the application or module, the product ignores the .xmi files.

However, a Java EE 5 or later module can exist within an application that includes pre-Java EE 5 files and uses the .xmi file name extension.

The `ibm-webservices-ext.xmi`, `ibm-webservices-bnd.xmi`, `ibm-webservicesclient-bnd.xmi`, `ibm-webservicesclient-ext.xmi`, and `ibm-portlet-ext.xmi` files continue to use the .xmi file extensions.

Web module

Specifies the name of a web module in the installed or deployed application.

URI

Specifies the location of the module that is relative to the root of the application (EAR file).

JSP enable class reloading

Specifies whether to enable class reloading when JSP files are updated.

A web container reloads JSP files only when the IBM extension `reloadEnabled` in the `jspAttributes` of the `ibm-web-ext.xmi` file is set to `true`.

Java Platform, Enterprise Edition 5 (Java EE 5) applications IBM extension files are in .xml file format. For applications versions earlier than Java EE 5, they are in the .xmi file format.

JSP reload interval in seconds

Specifies the number of seconds to scan the application file system for updated JSP files. The default is the value of the reloading interval attribute in the IBM extension (`META-INF/ibm-web-ext.xmi`) file of the web module.

To enable reloading, specify a value greater than zero (for example, 1 to 2147483647). The default reload interval is 5. To disable reloading, specify zero (0). The range is from 0 to 2147483647.

The reloading interval attribute takes effect only if class reloading is enabled.

Java EE 5 applications IBM extension files are in .xml file format. For applications versions earlier than Java EE 5, they are in the .xmi file format.

Sun Reference Implementation 1.2

Select this option to use the Sun Reference Implementation 1.2 JSF implementation.

If you change the JSF implementation that you are using for your application, you must delete any previously compiled JSP files. If you precompiled your application, you must recompile. If you did not precompile, but have already requested JSP files from this application, you must delete the JSP files from the temp directory of your profile.

You can set the JSF engine configuration parameter, `com.ibm.ws.jsf.JSF_IMPL_CHECK`, to `true` to automatically mark the JSP files to recompile at application startup.

MyFaces 2.0

Select this option to use the MyFaces JSF implementation. This is the default JSF implementation.

If you change the JSF implementation that you are using for your application, you must delete any previously compiled JSP files. If you precompiled your application, you must recompile. If you did not precompile, but have already requested JSP files from this application, you must delete the JSP files from the temp directory of your profile.

You can set the JSF engine configuration parameter, `com.ibm.ws.jsf.JSF_IMPL_CHECK`, to true to automatically mark the JSP files to recompile at application startup.

In a mixed-version cell, a V7 node uses MyFaces 1.2 if the MyFaces selection is toggled, while a V8 node uses MyFaces 2.0. For WebSphere Application Server versions before V7 (for example, V6.1 and earlier), this toggle is ineffective because JSF implementation switching was not supported before V7.

Context root for web modules settings

Use this page to specify the context root for web modules during or after installation of an application onto a WebSphere Application Server deployment target.

To view this administrative console page, click **Applications > Application Types > WebSphere enterprise applications > *application_name* > Context root for web modules**. This page is the same as the Context root for web modules page on the application installation and update wizards.

Web Module

Specifies the name of a web module in the application that you are installing or that you are viewing after installation.

URI

Specifies the location of the module relative to the root of the application EAR file.

Context Root

Specifies the context root of the web application (WAR).

A context root for each web module is defined in the application deployment descriptor during application assembly. Use this field to assign a different context root to a web module. 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`.

Initial parameters for servlets settings

Use this page to specify initial parameters that are passed to the init method of web module servlet filters. You can specify initial parameter values for servlets in web modules during or after installation of an application onto a WebSphere Application Server deployment target. The `<param-value>` values specified in `<init-param>` statements in the `web.xml` file of web modules are used by default.

To view this administrative console page, click **Applications > Application Types > WebSphere enterprise applications > *application_name* > Init parameters for servlets**. This page is the same as the Init parameters for servlets in each web module panel on the application installation and update wizards.

Module

Specifies the name of a module in the application that you are installing or that you are viewing after installation.

URI

Specifies the location of the module relative to the root of the application (EAR file).

Servlet

Specifies a unique name for the servlet within the application.

A *servlet* is a Java program that uses the Java Servlet Application Programming Interface (API). You must package servlets in a Web archive (WAR) file or web module for deployment to an 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.

Name

Specifies the name of the initial parameter passed to the init method of the web module servlet filter.

The following example servlet filter statement in a web.xml file specifies an initial parameter name of attribute:

```
<init-param>
  <param-name>attribute</param-name>
  <param-value>tests.Filter.DoFilter_Filter.SERVLET_MAPPED</param-value>
</init-param>
```

Value

Specifies the value assigned to an initial parameter passed to the init method of the web module servlet filter.

The following example servlet filter statement in a web.xml file specifies an initial parameter value of tests.Filter.DoFilter_Filter.SERVLET_MAPPED for the init parameter attribute:

```
<init-param>
  <param-name>attribute</param-name>
  <param-value>tests.Filter.DoFilter_Filter.SERVLET_MAPPED</param-value>
</init-param>
```

Description

Specifies information on the initial parameter.

Environment entries for client modules settings

Use this page to configure the environment entries of application client modules that are deployed as Java archive (JAR) files.

To view this administrative console page, click **Applications > Application Types > WebSphere enterprise applications > *application_name* > Environment entries for client modules**.

This page is the same as the Map environment entries for client modules page on the application installation and update wizards. To view the Map environment entries for client modules page in a wizard, you must select the **Deploy client modules** option on the Select installation options page.

Client module

Specifies the name of a client module.

URI

Specifies the location of the client module relative to the root of the application.

Name

Specifies the name of the environment entry that you are editing or viewing. The environment entry is the env-entry property in the client module.

Type

Specifies a data type for the environment entry defined by the env-entry property in the client module.

Description

Specifies information about the environment entry.

Value

Specifies an editable value for the environment entry. The value is defined by the env-entry property in the client module.

The lookup name is displayed in the **Value** column if the lookup name is configured in the application metadata. The lookup name is not editable. If you do not specify a value on this page, the lookup name is used for the value.

Environment entries for EJB modules settings

Use this page to configure the environment entries of Enterprise JavaBeans (EJB) modules such as entity, session, or message driven beans.

To view this administrative console page, click **Applications > Application Types > WebSphere enterprise applications > *application_name* > Environment entries for EJB modules**. This page is the same as the Map environment entries for EJB modules page on the application installation and update wizards.

Module

Specifies the name of an EJB module.

URI

Specifies the location of the EJB module relative to the root of the application.

Bean

Specifies the name of an enterprise bean that is contained by the module.

Name

Specifies the name of the environment entry that you are editing or viewing. The environment entry is the env-entry property in the EJB module.

Type

Specifies a data type for the environment entry defined by the env-entry property in the EJB module.

Description

Specifies information on the environment entry.

Value

Specifies an editable value for the environment entry defined by the env-entry property in the EJB module.

The lookup name is displayed in the **Value** column if the lookup name is configured in the application metadata. The lookup name is not editable. If you do not specify a value on this page, the lookup name is used for the value.

Environment entries for web modules settings

Use this page to configure the environment entries of Web modules such as servlets and JavaServer Pages (JSP) files.

To view this administrative console page, click **Applications > Application Types > WebSphere enterprise applications > *application_name* > Environment entries for web modules**. This page is the same as the Environment entries for web modules page on the application installation and update wizards.

Module

Specifies the name of a web module.

URI

Specifies the location of the module relative to the root of the application (EAR file).

Name

Specifies the name of the environment entry that you are editing or viewing. The environment entry is the env-entry property in the web module.

Type

Specifies a data type for the environment entry defined by the env-entry property in the web module.

Description

Specifies information on the environment entry.

Value

Specifies an editable value for the environment entry defined by the env-entry property in the web module.

The lookup name is displayed in the **Value** column if the lookup name is configured in the application metadata. The lookup name is not editable. If you do not specify a value on this page, the lookup name is used for the value.

Environment entries for application settings

Use this page to configure the environment entries of applications that are deployed as enterprise archive (EAR) files.

To view this administrative console page, click **Applications > Application Types > WebSphere enterprise applications > *application_name* > Environment entries for the application.**

This page is the same as the Map environment entries for application level page on the application installation and update wizards. To view this page, the application must define one or more environment entries.

Name

Specifies the name of the environment entry that you are editing or viewing. The environment entry is the env-entry property in the application.

Type

Specifies a data type for the environment entry defined by the env-entry property in the application.

Description

Specifies information about the environment entry.

Value

Specifies an editable value for the environment entry. The value is defined by the env-entry property in the application.

The lookup name is displayed in the **Value** column if the lookup name is configured in the application metadata. The lookup name is not editable. If you do not specify a value on this page, the lookup name is used for the value.

Resource environment references

Use this page to designate how the resource environment references of application modules map to remote resources, which are represented in the product as resource environment entries.

To view this administrative console page, click **Applications > Application Types > WebSphere enterprise applications > *application_name* > Resource environment references.**

Each row of the table depicts a resource environment reference within a specific module of your application. If you bound any references to resource environment entries during application assembly, you see the JNDI names of those resource environment entries in the applicable rows.

To set the mapping relationships between your resource environment references and resource environment entries:

1. Select a row. Be aware that if you check multiple rows on this page, the resource mapping target that you select in step 2 applies to all of those references.
2. Click **Browse** to select a resource environment entry from the new page that is displayed, the Available Resources page. The Available Resources page shows all resource environment entries that are available mapping targets for your application references.
3. Click **Apply**. The console displays the Resource environment references page again. In the rows that you previously selected, you now see the JNDI name of the new resource mapping target.
4. Repeat the previous steps as necessary.
5. Click **OK**. You now return to the general configuration page for your enterprise application.

Table column heading descriptions:

Select

Select the check boxes of the rows that you want to edit.

Module

The name of a module in the application.

EJB

The name of an enterprise bean that is accessed by the module.

URI

Specifies location of the module relative to the root of the application EAR file.

Reference binding

The name of a resource environment reference that is declared in the deployment descriptor of the application module. The reference corresponds to a resource that is bound as a resource environment entry into the JNDI name space of the application server.

JNDI name

The Java Naming and Directory Interface (JNDI) name of the resource environment entry that is the mapping target of the resource environment reference.

Data type String

Message destination reference settings

If your application uses message-driven beans, use this page to specify the Java Naming and Directory Interface (JNDI) name of the J2C administered object to bind the message destination reference to the message-driven beans. You must map each message destination reference that is defined in your application to an administered object.

To view this administrative console page, click **Applications > Application Types > WebSphere enterprise applications > *application_name* > Message destination references**. This page is the same as the **Bind message destination references to administered objects** page on the application installation and update wizards.

If the message destination reference is from an EJB 3.0 or later module, then the JNDI name is optional and the run time provides a container default value.

Attention: If multiple message destination references link to the same message destination, only one JNDI name is collected. When a message destination reference links to the same message destination as a message-driven bean and the destination JNDI name has been collected already, the destination JNDI name for the message destination reference is not collected.

To apply binding changes to multiple mappings:

1. In the list of mappings, select the **Select** check box beside each EJB module that you want mapped to a particular binding.
2. Expand **Apply Multiple Mappings**.
3. Complete one of the following steps:
 - Specify a message destination name.
 - Select a target resource JNDI name for a message destination.
4. Click **Apply**.
5. Click **OK** or **Next**.

Module

Specifies the name of the module that contains the bean.

Bean

Specifies name of a bean in the application.

URI

Specifies the location of the module relative to the root of the enterprise archive (EAR) file.

Message destination object

Specifies the message destination object.

Type

Specifies the type of object.

Target Resource JNDI Name

Specifies the Java Naming and Directory Interface (JNDI) name of the bean.

This is a data entry field. To change the JNDI name bound to this bean, type the new name in this field.

Select current backend ID settings

Use this page to select a backend identifier for container-managed persistence (CMP) beans that contain mappings for multiple backend databases.

This administrative console page is a step in the application installation and update wizards. To view this administrative console page, click **Applications > New Application > New Enterprise Application > application_path > Next > Detailed - Show all installation options and parameters > Next > Next** or **Continue > Step: Select current backend ID**. This page is displayed in the wizards if **Database type** is blank on the Provide options to perform the EJB Deploy page.

A backend can represent different database vendors, or simply alternative mappings and table qualifiers. If a Java archive (JAR) file for an enterprise bean defines CMP beans that contain mappings for multiple backend databases, you must select a current backend ID to be used when the module is installed on a deployment target. The backend ID determines the persister classes that get loaded at deployment.

Module

Specifies the name of the module that contains the bean.

URI

Specifies the location of the module relative to the root of the application EAR file.

Current backend ID

Specifies the current backend ID to be used when the module is installed on a deployment target.

Provide JNDI names for JCA objects settings

Use this page to configure Java Naming and Directory Interface (JNDI) name values for J2C objects (J2CConnectionFactory, J2CActivationSpec, and J2CAdminObject) in your application or modules. If your application contains an embedded resource archive (RAR) file, specify the name and JNDI name of each JCA connection factory, administered object, and activation specification.

This administrative console page is a step in the application installation and update wizards. To view this administrative console page, click **Applications > New Application > New Enterprise Application > application_path > Next > Detailed - Show all installation options and parameters > Next > Next or Continue > Step: Provide JNDI names for JCA objects**.

Connector module

Specifies the name of a connector module of the RAR file.

URI

Specifies the location of the module that is relative to the root of the RAR file.

Object identifier

Specifies the name of the J2C object. The object can be a JCA connection factory, administered object, or activation specification.

Bindings

Specifies the name and Java Naming and Directory Interface (JNDI) name of the J2C object.

These are data entry fields. To change the name or JNDI name bound to this object, type the new names in the fields.

Correct use of the system identity

Use this page to manage the system identity properties for the Enterprise JavaBeans (EJB) method in your application.

This administrative console page is displayed during the application deployment process. To access the administrative console, click **Application > New application > New Enterprise Application**. The is displayed as *Correct use of System Identity* in the application deployment steps.

To use this page, complete the following steps:

1. Select an application that supports security and click **Next**.
2. Select **Detailed - Show all installation options and parameters** and click **Next**.
3. Select the **Correct use of system identity** step.

Bean

A component that implements a business task or business entity and resides in an EJB container. Entity beans, session beans, and message-driven beans are all enterprise beans.

Module

In Java EE programming, a software unit that consists of one or more components of the same container type and one deployment descriptor of that type. Examples include EJB, Web, and application client modules.

URI

A Uniform Resource Identifier (URI) is a unique address that is used to identify content on the Web, such as a page of text, a video or sound clip, a still or animated image, or a program.

Method signature

The combination of a name of a method along with the number and types of the parameters and their order.

Role

Specifies the RunAs role that is used for this EJB method.

Username

Specifies the user name that is assigned to the RunAs role for this EJB method.

The user name is used in conjunction with the RunAs role that you select for the Role.

Requirements for setting data access isolation levels

This article discusses the criteria and effects of setting isolation levels for data access components that comprise Enterprise JavaBeans (EJB) 2.x and later modules.

In an EJB 1.1 module, you can set the isolation level at the method level or bean level. This capability also applies to container-managed persistence (CMP) 1.1 beans that you assemble into *EJB 2.x modules*. WebSphere Application Server permits the deployment descriptor of a CMP bean to declare the version level of 1.1, regardless of the overall module version.

However, the ability to set isolation level at the method or bean level does **not** apply to other enterprise beans within an EJB 2.x module, including *CMP 2.x beans*. WebSphere Application Server Version 5.0 removed this capability from EJB 2.0 modules to deliver an architecture that ultimately provides more efficient connection use.

Consequently, later versions of the product enforce the following restrictions on declaring isolation level for CMP 2.x beans—as well as session beans, message-driven beans, and bean managed persistence (BMP) beans that you assemble into EJB 2.x modules:

- You cannot specify isolation level on the EJB method level or bean level.
- If you configure a JDBC application, a bean-managed persistence (BMP) bean, or a servlet to participate in global transactions, any connection that is shared cannot accept a user-specified isolation level. WebSphere Application Server can only set a user-specified isolation level on a connection that is not shared within a global transaction. *Generally, you want to refrain from specifying isolation levels on shareable connections.*

The configuration for the isolation level is determined by the type of bean that is used by the component:

Isolation level on connections used by 2.x CMP beans

In a EJB 2.x module, when a CMP 2.x bean uses a new data source to access a backend database, the isolation level is determined by the WebSphere Application Server run time, based on the type of access intent assigned to the bean or the calling method. Other non-CMP connection users can access this same data source and also use the access intent and application profile support to manage their concurrency control.

Connections used by other 2.x enterprise beans and other non-CMP components

For all other JDBC connection instances (connections other than those used by CMP beans), you can specify an isolation level on the data source resource reference. For shareable connections that run in global transactions, this method is the only way to set the *isolationLevel* for connections. Trying to directly set the isolation level through the *setTransactionIsolation()* method

on a shareable connection that runs in a global transaction is not allowed. To use a different isolation level on connections, you must provide a different resource reference. Set these defaults through your assembly tool.

Each resource reference associates with one isolation level. When your application uses this resource reference Java Naming and Directory Interface (JNDI) name to look up a data source, every connection returned from this data source using this resource reference has the same isolation level.

Components needing to use shareable connections with multiple isolation levels can create multiple resource references, giving them different JNDI names, and have their code look up the appropriate data source for the isolation level they need. In this way, you use separate connections with the different isolation levels enabled on them.

It is possible to map these multiple resource references to the same configured data source. The connections still come from the same underlying pool, however; the connection manager does not allow sharing of connections requested by resource references with different isolation levels. Consider the following scenario:

- A data source is bound to two resource references: *jdbc/RRResRef* and *jdbc/RResRef*.
- RRResRef has the *RepeatableRead* isolation level defined. RResRef has the *ReadCommitted* isolation level defined.

If your application wants to update the tables or a BMP bean updates some attributes, it can use the *jdbc/RRResRef* JNDI name to look up the data source instance. All connections returned from the data source instance have a RepeatableRead isolation level. If the application wants to perform a query for read only, then it is better to use the *jdbc/RResRef* JNDI name to look up the data source.

If you do not specify the isolation level:

The product does not require you to set the isolation level on a data source resource reference for a non-CMP application module. If you do not specify isolation level on the resource reference, or if you specify TRANSACTION_NONE, the WebSphere Application Server run time uses a default isolation level for the data source. Application Server uses a default setting based on the JDBC driver.

For most drivers, WebSphere Application Server uses an isolation level default of TRANSACTION_REPEATABLE_READ. For Oracle drivers, however, Application Server uses an isolation level of TRANSACTION_READ_COMMITTED. Use the following table for quick reference:

Database:	DB2	Oracle	Sybase	Informix®	Apache Derby	SQL Server
Default isolation level: <i>(for connections used by non-CMP entities)</i>	RR	RC	RR	RR	RR	RR

- **Note:** These same default isolation levels are used in cases of direct JNDI lookups of a data source.
- RR = JDBC Repeatable read (TRANSACTION_REPEATABLE_READ)
- RC = JDBC Read committed (TRANSACTION_READ_COMMITTED)

To customize the default isolation level, you can use the `webSphereDefaultIsolationLevel` custom property for the data source. In most cases you should define the isolation level in the deployment descriptor when you package the EAR file, but in certain situations you might need to customize

the default isolation level. This property will have no effect if any of the above options are used, and this custom property is provided for those situations in which there is no other means of setting the isolation level.

Use the following values for `webSphereDefaultIsolationLevel` custom property:

Possible values	JDBC isolation level	DB2 isolation level
8	TRANSACTION_SERIALIZABLE	Repeatable Read (RR)
4 (default)	TRANSACTION_REPEATABLE_READ	Read Stability (RS)
2	TRANSACTION_READ_COMMITTED	Cursor Stability (CS)
1	TRANSACTION_READ_UNCOMMITTED	Uncommitted Read (UR)

To define this custom property for a data source:

1. Click **Resources > JDBC provider > JDBC_provider**.
2. Click **Data sources** in the Additional Properties section.
3. Click the name of the data source.
4. Click **Custom properties**.
5. Create the `webSphereDefaultIsolationLevel` custom property.
 - a. Click **New**.
 - b. Enter `webSphereDefaultIsolationLevel` for the name field.
 - c. Enter one of the possible values in the value field.

Application Server sets the isolation level by prioritizing the available settings. Application Server will set the isolation level based on the values for the following, in this order:

1. Resource reference isolation level
2. Isolation level that is specified by the access intent policy
3. Custom property that configures an isolation level
4. Application Server's default setting.

Metadata for module settings

Use this page to instruct a Java Platform, Enterprise Edition (Java EE) enterprise bean (EJB) deployment descriptor, web module deployment descriptor, or JCA resource adapter archive (RAR) module to ignore annotations that specify deployment information.

To view this administrative console page, click **Applications > Application Types > WebSphere enterprise applications > application_name > Metadata for modules**. This page is the same as the Metadata for modules page on the application installation and update wizards.

If your application contains Java EE 5 or later modules, you can select to lock the deployment descriptor of one or more of the modules on the Metadata for modules page. If you select a **metadata-complete attribute** check box (set the `metadata-complete` attribute to `true`) and lock deployment descriptors, the product writes the complete module deployment descriptor, including deployment information from annotations, to XML format.

Annotations are a standard mechanism of adding metadata to Java classes. You can use metadata to simplify development and deployment of Java EE 5 or later artifacts. Prior to the introduction of Java language annotations, deployment descriptors were the standard mechanism used by Java EE components. These deployment descriptors were mapped to XML format, which facilitated their persistence. If you select to lock deployment descriptors, the product merges Java EE annotation-based metadata with the XML-based existing deployment descriptor metadata and persists the result.

When applications contain a large number of Java classes, the deployment processing time for the annotations can increase. To minimize the performance impact, you can use one of the following methods:

- Determine whether the module needs to use Java EE 5 or 6. If the module does not need to use Java EE 5 or 6, the annotations within the Java classes are not scanned.
- Use the “metadata-complete attribute” in the module descriptor if the module uses Java EE 5 or later and it does not contain any annotations. This attribute disables the annotations processing for the module, but Java EE 5 or later modules might still be placed in the descriptor file. If you are migrating your application, but you are not adding annotations, consider using this attribute value.
- Restructure the application to place the utility Java archive (JAR) files into shared libraries if those JAR files do not have annotation information. Consider this method if you cannot set the “metadata-complete attribute.”
- Move the JAR files in the WEB-INF/lib directory to the root directory of the enterprise archive (EAR) file. Nested archives, such as a JAR file that is within a web application archive (WAR) that is within an EAR file, are very cumbersome to search through because of the multiple levels of compression.

Module

Specifies the name of a module in the installed (or deployed) application.

Data type String

URI

Specifies the location of the module relative to the root of the EAR file.

Data type String

metadata-complete attribute

Specifies whether to write the complete module deployment descriptor, including deployment information from annotations, to extensible markup language (XML) format.

By default, a **metadata-complete attribute** check box is not selected and the product does not write out annotation data to a module deployment descriptor.

If your modules do not have a metadata-complete attribute or the metadata-complete attribute is set to false, you can select a check box and instruct the product to write out annotation data to a module deployment descriptor.

Note: If your Java EE 5 or later application uses annotations and a shared library, do not select **metadata-complete attribute**. When your application uses annotations and a shared library, setting the metadata-complete attribute to true causes the product to incorrectly represent an @EJB annotation in the deployment descriptor as <ejb-ref> rather than <ejb-local-ref>. For web modules, setting the metadata-complete attribute to true might cause InjectionException errors. If you must select **metadata-complete attribute** (set the metadata-complete attribute to true), avoid errors by not using a shared library, by placing the shared library in either the classes or lib directory of the application server, or by fully specifying the metadata in the deployment descriptors.

After you select a check box, you cannot deselect (clear) the check box and the module is no longer shown in the list of modules on this page. If you select all the check boxes, the link to this page is no longer shown on the enterprise application settings page.

Data type Boolean
Default false (deselected)

Provide options to perform the web services deployment settings

Use this page to specify options for web services deployment.

This administrative console page is a step in the application installation and update wizards.

To view this page, you must select **Deploy web services** on the **Select installation options** page.

To view this administrative console page, complete the following steps:

1. Click **Applications > New application > *application_path***.
2. Select the option to **Show all installation options and parameters**.
3. Click **Next** to get to the **Step: Select installation options** page.
4. Select **Deploy web service**.
5. Click **Next** to get to the **Step: Provide options to perform the web services deployment** page.

You can specify the web services deployment options on this page only when installing or updating an application that uses web services.

The `wsdeploy` command is supported by Java API for XML-based RPC (JAX-RPC) applications. The Java API for XML-Based Web Services (JAX-WS) programming model that is implemented by the application server does not support the `wsdeploy` command. If your web services application contains only JAX-WS endpoints, you do not need to run the `wsdeploy` command, as this command is used to process only JAX-RPC endpoints.

The options that you specify set parameter values for the `wsdeploy` command. The `wsdeploy` command adds product-specific deployment classes to a web services-compatible enterprise archive (EAR) file or an application client Java archive (JAR) file. These classes include:

- Stubs
- Serializers and deserializers
- Implementations of service interfaces

The `wsdeploy` command is run during installation after you click **Finish** on the **Summary** page of the wizard.

Deploy web services option - Classpath

Specifies entries to add to the CLASSPATH when the generated classes are compiled.

To specify the class paths of multiple entries, you need to separate the entries with a semicolon on Windows platforms and on Linux, Unix, and z/OS platforms, you need to use a colon to separate the entries. This is the same separator that is used with the CLASSPATH environment variable.

This option is the same as the `wsdeploy` command parameter `-cp class_path`.

Data type	String
Default	null

Deploy web services option - Extension Directories

Specifies a directory that contains zipped or Java archive (JAR) files. All zipped and JAR files in this directory are added to the CLASSPATH used to compile the generated files.

This option is the same as the `wsdeploy` command parameter `-jardir directory`.

Data type	String
Default	null

Display module build ID settings

Use this page to view the build identifier of a module in a Java Platform, Enterprise Edition (Java EE) enterprise archive (EAR file). The build identifier for a module is shown if the MANIFEST.MF file of a module or application specifies a build identifier.

To view this administrative console page, click **Applications > Application Types > WebSphere enterprise applications > *application_name* > Display module build IDs**. This page is the same as the Display module build IDs page on the application installation and update wizards.

Module

Specifies the name of a module in the installed (or deployed) application.

Data type String

URI

Specifies the location of the module relative to the root of the application EAR file.

Data type String

Build ID

Specifies the build identifier for a module if the MANIFEST.MF file specifies a build identifier.

You cannot modify the build ID on this page because this field is read-only.

Data type String

Task overview: Assembling applications using remote request dispatcher

Remote request dispatcher (RRD) is a pluggable extension to the web container which allows application frameworks, servlets and JavaServer Pages (JSP) to include content from outside the currently executing resource's Java Virtual Machine (JVM) as part of the response sent to the client.

Before you begin

You must have WebSphere Application Server, Network Deployment installed to use remote request dispatcher function. You should also familiarize yourself the limitations of remote request dispatcher. See article, Remote request dispatcher considerations for details.

Procedure

1. Install enterprise application files with the console.
2. Configure the sending of include requests between the application and remote resources.
 - Configure web applications to dispatch remote includes.
 - Configure web applications to service remote includes.
3. Optional: Modify your application to locate resources located in two different contexts using the servlet programming model.

The Servlet Programming Model for including resources remotely does not require you to use any non-Java Platform, Enterprise Edition (Java EE) Servlet Application Programming Interfaces (APIs). The remote request dispatcher (RRD) component follows the same rules to obtain a ServletContext and a remote resource. By using JavaServer Pages standard tag library (JSTL), your application is further removed from obtaining a ServletContext object or RequestDispatcher that is required in the

framework example in the following step because the JSTL custom tag does this implicitly. Study the following example of a sample JavaServer Pages application to learn how to locate resources that are in two different contexts, investments and banking.

```
<HEAD>
<%@ page
language="java"
contentType="text/html; charset=ISO-8859-1"
pageEncoding="ISO-8059-1"
isELIgnored="false"
%>
<%@ taglib uri="http://java.sun.com/jsp/jstl/core" prefix="c" %>
</HEAD>
<BODY>

<%--
```

Programming example using JavaServer Pages and JavaSever Pages Standard Tag Library (JSTL). JSTL provides a custom tag to import contents (in servlet and JSP terms include) in the scope of the same request from outside of the current web module context by specifying a context parameter.

JSTL restriction: The web module that is imported must run inside of the same JVM as the calling resource if imported URL is not fully qualified.

RRD extends this functionality by permitting the web module to be located within the scope of the current WebSphere Application Server core group versus the scope of the JVM.

```
--%>

<hr size="5"/>
<%-- Include resource investmentSummary.jsp located in the
web application with context root of /investments. --%>

<c:import url="investmentSummary.jsp" context="/investments"/>

<hr size="5"/>
<%-- Include resource accountSummary.jsp located in the
web application with context root of /banking. --%>

<c:import url="accountSummary.jsp" context="/banking"/>

<hr size="5"/>
</BODY>
</HTML>
```

4. Optional: Modify your application to locate resources located in two different contexts using the framework programming model.

The Framework Programming Model for including resources remotely does not require you to use any non- Java Platform, Enterprise Edition (Java EE) Servlet Application Programming Interfaces (APIs). When a request is initiated for a ServletContext name that is not presently running inside of the current web container, the remote request dispatcher (RRD) component returns a ServletContext object that can locate a resource that exists anywhere inside a WebSphere Application Sever WebSphere Application Server, Network Deployment environment provided that the resource exists and RRD is enabled for that ServletContext object. Study the following sample framework snippet that demonstrates how to locate resources located in two different contexts, investments and banking.

```
/*
Programming example using a generic framework.
Servlet Specification provides an API to obtain
a servlet context in the scope of the same request
different from the current web module context by
specifying a context parameter.

Servlet Specification restriction: The web module that obtain
must run inside of the same JVM as the calling resource.

RRD extends this functionality by permitting the web module to be located
within the scope of the current WebSphere Application Server core group
versus the scope of the JVM.

*/

protected void frameworkCall (ServletContext context, HttpServletRequest request, HttpServletResponse response)
    throws ServletException, IOException{

    PrintWriter writer = response.getWriter();

    writer.write("<HTML>");
```

```

writer.write("<HEAD>");
writer.write("</HEAD>");
writer.write("<BODY>");
writer.write("<hr size=\\"5/\>");

//Include resource investmentSummary.jsp located in web application
//with context root of /investments.
RequestDispatcher rd = getRequestDispatcher ( context, "/investments", "/investmentSummary.jsp");
rd.include(request, response);

writer.write("<hr size=\\"5/\>");

//Include resource accountSummary.jsp located in web application
//with context root of /banking.
rd = getRequestDispatcher ( context, "/banking", "/accountSummary.jsp");
rd.include(request, response);

writer.write("</BODY>");
writer.write("</HTML>");
}
private RequestDispatcher getRequestDispatcher (ServletContext context, String contextName, String resource) {
return context.getContext(contextName).getRequestDispatcher(resource);
}

```

Results

After enabling at least one enterprise application to dispatch remote includes and at least one enterprise application to service remote includes, RRD is now enabled.

What to do next

Restart the modified applications if already installed or start newly installed applications to enable RRD on each application.

Installing enterprise modules with JSR-88

You can install Java Platform, Enterprise Edition (Java EE) modules on an application server provided by a WebSphere Application Server product using the Java EE Application Deployment API specification (JSR-88).

Before you begin

Note: Application installation using the Java EE Application Deployment API specification (JSR-88) has been deprecated in WebSphere Application Server Version 8.0. Use another option to deploy applications to a server. The closest option to using the Java EE Deployment API is using Java Management Extensions (JMX) MBean programming. For information on deployment options, see "Ways to install enterprise applications or modules."

JSR-88 defines standard application programming interfaces (APIs) to enable deployment of Java EE applications and stand-alone modules to Java EE product platforms. The Java EE Application Deployment specification Version 1.1 is available at <http://java.sun.com/j2ee/tools/deployment/reference/docs/index.html> as part of the Java 2 Platform, Enterprise Edition (J2EE) 1.4 Application Server Developer Release.

Read about JSR-88 and APIs used to manage applications at <http://java.sun.com/j2ee/tools/deployment/>.

About this task

JSR-88 defines a contract between a tool provider and a platform that enables tools from multiple vendors to configure, deploy and manage applications on any Java EE product platform. The tool provider typically supplies software tools and an integrated development environment (IDE) for developing and assembly of Java EE application modules. The Java EE platform provides application management functions that deploy, undeploy, start, stop, and otherwise manage Java EE applications.

WebSphere Application Server is a Java EE specification-compliant platform that implements the JSR-88 APIs. Complete the following steps to deploy (install) Java EE modules on an application server provided by the WebSphere Application Server platform.

Procedure

1. Code a Java program that can access the JSR-88 DeploymentManager class for the product.
 - a. Write code that finds the JAR manifest attribute J2EE-DeploymentFactory-Implementation-Class. Under JSR-88, your code finds the DeploymentFactory using the JAR manifest attribute J2EE-DeploymentFactory-Implementation-Class. The following product application management JAR files contain this attribute and provide support.

Table 27. JAR files that contain the manifest attribute. Enable your code to find the DeploymentFactory using the JAR manifest attribute.

Environment	JAR file containing the manifest attribute
Application server	<i>app_server_root</i> /plugins/com.ibm.ws.admin.services.jar
Application client	<i>app_client_root</i> /plugins/com.ibm.ws.j2ee.client.jar
Thin application client	<i>app_client_root</i> /runtimes/com.ibm.ws.admin.client_8.0.0.jar

After your code finds the DeploymentFactory, the deployment tool can create an instance of the WebSphere DeploymentFactory and register the instance with its DeploymentFactoryManager.

Example code for the application server environment follows. The example code requires that you use the development kit shipped with the product or use the pluggable client for deployment of stand-alone modules. See *WebSphere Application Server detailed system requirements* at <http://www.ibm.com/support/docview.wss?rs=180&uid=swg27006921> for information on supported development kits.

```
import javax.enterprise.deploy.shared.factories.DeploymentFactoryManager;
import javax.enterprise.deploy.spi.DeploymentManager;
import javax.enterprise.deploy.spi.factories.DeploymentFactory;
import java.util.jar.JarFile;
import java.util.jar.Manifest;

// Get the DeploymentFactory implementation class from the MANIFEST.MF file.
File jsr88Jar = new File(wasHome + "/plugins/com.ibm.ws.admin.services.jar");
JarFile jarFile = new JarFile(jsr88Jar);
Manifest manifest = jarFile.getManifest();
Attributes attributes = manifest.getMainAttributes();
String key = "J2EE-DeploymentFactory-Implementation-Class";
String className = attributes.getValue(key);
// Get an instance of the DeploymentFactoryManager
DeploymentFactoryManager dfm = DeploymentFactoryManager.getInstance();

// Create an instance of the WebSphere Application Server DeploymentFactory.
Class deploymentFactory = Class.forName(className);
DeploymentFactory deploymentFactoryInstance =
    (DeploymentFactory) deploymentFactory.newInstance();

// Register the DeploymentFactory instance with the DeploymentFactoryManager.
dfm.registerDeploymentFactory(deploymentFactoryInstance);

// Provide WebSphere Application Server URI, user ID, and password.
// For more information, see the step that follows.
wsDM = dfm.getDeploymentManager(
    "deployer:WebSphere:myserver:8880", null, null);
```

- b. Write code that accesses the DeploymentManager instance for the product.

The product URI for deployment has the following format:

```
"deployer:WebSphere:host:port"
```

The example in the previous step, "deployer:WebSphere:myserver:8880", tries to connect to host myserver at port 8880 using the SOAP connector, which is the default.

You can specify an Internet Protocol Version 6 (IPv6) address for the *host* element in the URI for deployment. Enclose the IPv6 address in square brackets ([]); for example:

```
"deployer:WebSphere:[IPv6_address]:port"
```

Also, you can add an optional parameter, *connectorType*, to the URI for deployment. For example, to use the RMI connector to access myserver, code the URI as follows:

```
"deployer:WebSphere:myserver:2809?connectorType=RMI"
```

2. Optional: Code a Java program that can customize or deploy Java EE applications or modules using the JSR-88 support provided by the product.
3. Start the deployed Java EE applications or stand-alone Java EE modules using the JSR-88 API used to start applications or modules.

What to do next

Test the deployed applications or modules. For example, point a web browser at the URL for a deployed application and examine the performance of the application. If necessary, update the application.

Customizing modules using DConfigBeans

You can configure Java Platform, Enterprise Edition (Java EE) applications or stand-alone modules during deployment using the DConfigBean class in the Java EE Application Deployment API specification (JSR-88).

Before you begin

Note: Application installation using the Java EE Application Deployment API specification (JSR-88) has been deprecated in WebSphere Application Server Version 8.0. Use another option to deploy applications to a server. The closest option to using the Java EE Deployment API is using Java Management Extensions (JMX) MBean programming. For information on deployment options, see "Ways to install enterprise applications or modules."

This topic assumes that you are deploying (installing) Java EE modules on an application server provided by the product using the WebSphere Application Server support for JSR-88.

Read about the JSR-88 specification and using the DConfigBean class at <http://java.sun.com/j2ee/tools/deployment/>.

About this task

The DConfigBean class in JSR-88 provides JavaBeans-based support for platform-specific configuration of J2EE applications and modules during deployment. Your code can inspect DConfigBean instances to get platform-specific configuration attributes. The DConfigBean instances provided by WebSphere Application Server contain a single attribute which has an array of java.util.Map objects. The map entries contain configuration attributes, for which your code can get and set values.

Procedure

1. Write code that installs Java EE modules on an application server using JSR-88.
2. Write code that accesses DConfigBeans generated by the product during JSR-88 deployment. You (or a deployer) can then customize the accessed DConfigBeans instances.

The following pseudocode shows how a Java EE tool provider can get DConfigBean instance attributes generated by the product during JSR-88 deployment and set values for the attributes.

```
import javax.enterprise.deploy.model.*;
import javax.enterprise.deploy.spi.*;
{
DeploymentConfiguration dConfig = __; // Get from DeploymentManager
DDBeanRoot ddRoot = __; // Provided by J2EE tool

// Obtain root bean.
DConfigBeanRoot dcRoot = dConfig.getDConfigBeanRoot(dr);

// Configure DConfigBean.
configureDCBean (dcRoot);
}

// Get children from DConfigBeanRoot and configure each child.
method configureDCBean (DConfigBean dcBean)
{
// Get DConfigBean attributes for a given archive.
BeanInfo bInfo = Introspector.getBeanInfo(dcBean.getClass());
```

```

IndexedPropertyDescriptor ipDesc =
    (IndexedPropertyDescriptor)bInfo.getPropertyDescriptors()[0];

// Get the 0th map.
int index = 0;
Map map = (Map)
    ipDesc.getIndexedReadMethod().invoke
        (dcBean, new Object[]{new Integer(index)});

while (map != null)
{
    // Iterate over the map and set values for attributes.

    // Set the map back into the DCBean.
    ipDesc.getIndexedWriteMethod().invoke
        (dcBean, new Object[]{new Integer(index), map});

    // Get the next entry in the indexed property
    map = (Map)
        ipDesc.getIndexedReadMethod().invoke
            (dcBean, new Object[]{new Integer(++index)});
}
}

```

Configuring enterprise application files

You can change the configuration of a Java Platform, Enterprise Edition (Java EE) application or module deployed on a server.

Before you begin

You can change the contents and deployment descriptors of an application or module before deployment, such as in an assembly tool. However, it is assumed that the module is already deployed on a server.

About this task

Changing an application or module configuration consists of one or more of the following:

- Changing the settings of the application or module.
- Removing a file from an application or module.
- Updating the application or its modules.

This topic describes how to change the settings of an application or module using the administrative console.

Attention: If an application is running, changing an application setting causes the application to restart. On stand-alone servers, the application restarts after you save the change. On multiple-server products, the application restarts after you save the change and files synchronize on the node where the application is installed. To control when synchronization occurs on multiple-server products, deselect **Synchronize changes with nodes** on the Console preferences page.

Procedure

- View current settings of the application or module.

Click **Applications > Application Types > WebSphere enterprise applications > application_name** to access the enterprise application settings page.

Many application or module settings are available on other console pages that you can access by clicking links on the settings page for the enterprise application. For detailed information on the settings and allowed values, examine the online help for the console pages. When you installed the application or module, you specified most of the settings values.
- Map each module of your application to a target server.

Specify the application servers, clusters of application servers, or web servers onto which to install modules of your application.
- Change how quickly your application starts compared to other applications or to the server.
- Configure the use of binary files.
- Change how your application or web modules use class loaders.

- Map a virtual host for each web module of your application.
- Change application bindings or other settings of the application or module.
 1. Click **Applications > Application Types > WebSphere enterprise applications > *application_name* > *property_or_item_name***. From the enterprise application settings page, you can access console pages for further configuring of the application or module.
 2. Change the values for settings as needed, and click **OK**.
- Optional: Configure the application so it does not start automatically when the server starts. By default, an installed application starts when the server on which the application resides starts. You can configure the target mapping for the application so the application does not start automatically when the server starts. To start the application, you must then start it manually.
- If the installed application or module uses a resource adapter archive (RAR file), ensure that the **Classpath** setting for the RAR file enables the RAR file to find the classes and resources that it needs. Examine the **Classpath** setting on the console Resource adapter settings page.

Results

The application or module configuration is changed. The application or standalone web module is restarted so the changes take effect.

If you updated module metadata while the application was running, restarting the application might not be sufficient for the changes to take effect. For example, if you changed descriptors in running Java EE 6 applications that use annotations, you must reinstall the application. If you changed classes that introduce, remove, or alter class hierarchies within an application, and those changes impact annotated classes, you also must reinstall the application.

What to do next

If the application or module is deployed on a cluster and you have no more configuration changes to make, click **Rollout Update** on the Enterprise applications page to propagate the changed configuration on all cluster members of the cluster on which the application or module is deployed. **Rollout Update** sequentially updates the configuration on the nodes that contain cluster members.

Save changes to your administrative configuration.

On multiple-server products, the application binaries are transferred to nodes when the configuration changes on the deployment manager synchronize with configurations for individual nodes on which the application will run.

Application bindings

Before an application that is installed on an application server can start, all enterprise bean (EJB) references and resource references defined in the application must be bound to the actual artifacts (enterprise beans or resources) defined in the application server.

When defining bindings, you specify Java Naming and Directory Interface (JNDI) names for the referenceable and referenced artifacts in an application. The `jndiName` values specified for artifacts must be qualified lookup names. An example referenceable artifact is an EJB defined in an application. An example referenced artifact is an EJB or a resource reference used by the application.

Binding definitions are stored in the `ibm-xxx-bnd.xml` or `ibm-xxx-bnd.xmi` files of an application. Version 7.0 or later binding definitions support files with the suffix of XML for EJB 3.x and Web 2.5 and later modules. Modules earlier than Java EE 5 continue to use binding definition files with the suffix of XMI as in previous versions of WebSphere Application Server. The `xxx` can be `ejb-jar`, `web`, `application` or `application-client`.

This topic provides the following information about bindings:

- “Times when bindings can be defined”
- “Required bindings”
- “Application resource conflicts” on page 289

Times when bindings can be defined

You can define bindings at the following times:

- During application development

An application developer can create binding definitions in `ibm-xxx-bnd.xml` files for EJB 3.x and Web 2.5 and later modules and in `ibm-xxx-bnd.xmi` files for pre-Java EE 5 modules. The application developer can create the files using a tool such as an IBM Rational developer tool or, for EJB 3.x or Web 2.5 and later modules, using an XML editor or text editor. The developer then gives an enterprise application (.ear file) complete with bindings to an application assembler or deployer. When assembling the application, the assembler does not modify the bindings. Similarly, when installing the application onto a server supported by WebSphere Application Server, the deployer does not modify or override the bindings or generate default bindings unless changes to the bindings are necessary for successful deployment of the application.

- During application assembly

An application assembler can define bindings in annotations or deployment descriptors of an application. Java EE 5 or later modules contain annotations in the source code. To declare an annotation, an application assembler precedes a keyword with an @ character ("at" sign). Bindings for pre-Java EE 5 modules are specified in the **WebSphere Bindings** section of a deployment descriptor editor. Modifying the deployment descriptors might change the binding definitions in the `ibm-xxx-bnd.xmi` files created when developing an application. After defining the bindings, the assembler gives the application to a deployer. When installing the application onto a server supported by WebSphere Application Server, the deployer does not modify or override the bindings or generate default bindings unless changes to the bindings are necessary to deploy the application.

- During application installation

An application deployer or server administrator can modify the bindings when installing the application onto a server supported by WebSphere Application Server using the administrative console. New binding definitions can be specified on the installation wizard pages.

Also, a deployer or administrator can select to generate default bindings during application installation. Selecting **Generate default bindings** during application installation instructs the product to fill in incomplete bindings in the application with default values. Existing bindings are not changed.

Restriction: You cannot define or override bindings during application installation for application clients. You must define bindings for application client modules during assembly and store the bindings in the `ibm-application-client-bnd.xmi` file.

- During configuration of an installed application

After an application is installed onto a server supported by WebSphere Application Server, an application deployer or server administrator can modify the bindings by changing values in administrative console pages such as those accessed from the settings page for the enterprise application.

Required bindings

Before an application can be successfully deployed, bindings must be defined for references to the following artifacts:

EJB JNDI names

For each EJB 2.1 or earlier enterprise bean (EJB), you must specify a JNDI name. The name is used to bind an entry in the global JNDI name space for the EJB home object. An example JNDI

name for a *Product* EJB in a *Store* application might be `store/ejb/Product`. The binding definition is stored in the `META-INF/ibm-ejb-jar-bnd.xml` file.

If a deployer chooses to generate default bindings when installing the application, the installation wizard assigns EJB JNDI names having the form *prefix/EJB_name* to incomplete bindings. The default prefix is `ejb`, but can be overridden. The *EJB_name* is as specified in the deployment descriptor `<ejb-name>` tag.

During and after application installation, EJB JNDI names can be specified on the Provide JNDI names for beans page. After installation, click **Applications > Application Types > WebSphere enterprise applications > *application_name* > EJB JNDI names** in the administrative console.

You do not need to specify JNDI binding names for each of the EJB 3.x home or business interfaces on your enterprise beans because the EJB container assigns default bindings.

Data sources for entity beans

Entity beans such as container-managed persistence (CMP) beans store persistent data in data stores. With CMP beans, an EJB container manages the persistent state of the beans. You specify which data store a bean uses by binding an EJB module or an individual enterprise bean to a data source. Binding an EJB module to a data source causes all entity beans in that module to use the same data source for persistence.

An example JNDI name for a *Store* data source in a *Store* application might be `store/jdbc/store`. For modules earlier than Java EE 5, the binding definition is stored in IBM binding files such as `ibm-ejb-jar-bnd.xml`. A deployer can also specify whether authentication is handled at the container or application level.

WebSphere Application Server Version 8.0 supports CMP beans in EJB 2.x or 1.x modules. Version 8.0 does not support CMP beans in EJB 3.0 modules.

If a deployer chooses to generate default bindings when installing the application, the installation wizard generates the following for incomplete bindings:

- For EJB 2.x `.jar` files, connection factory bindings based on the JNDI name and authorization information specified
- For EJB 1.1 `.jar` files, data source bindings based on the JNDI name, data source user name and password specified

The generated bindings provide default connection factory settings for each EJB 2.x `.jar` file and default data source settings for each EJB 1.1 `.jar` file in the application being installed. No bean-level connection factory bindings or data source bindings are generated unless they are specified in the custom strategy rule supplied during default binding generation.

During and after application installation, you can map data sources to 2.x entity beans on the 2.x CMP bean data sources page and on the 2.x entity bean data sources page. After installation, click **Applications > Application Types > WebSphere enterprise applications > *application_name*** in the administrative console, then select **2.x CMP bean data sources** or **2.x entity bean data sources**. You can map data sources to 1.x entity beans on the Map data sources for all 1.x CMP beans page and on the Provide default data source mapping for modules containing 1.x entity beans page. After installation, access console pages like those for 2.x CMP beans, except click links for 1.x CMP beans.

Backend ID for EJB modules

If an EJB `.jar` file that defines CMP beans contains mappings for multiple backend databases, specify the appropriate backend ID that determines which persister classes are loaded at run time.

Specify the backend ID during application installation. You cannot select a backend ID after the application is installed onto a server.

To enable backend IDs for individual EJB modules:

1. During application installation, select **Deploy enterprise beans** on the Select installation options page. Selecting **Deploy enterprise beans** enables you to access the Provide options to perform the EJB Deploy page.

2. On the Provide options to perform the EJB Deploy page, set the database type to "" (null).

During application installation, if you select **Deploy enterprise beans** on the Select installation options page and specify a database type for the EJB deployment tool on the Provide options to perform the EJB Deploy page, previously defined backend IDs for all of the EJB modules are overwritten by the chosen database type.

The default database type is DB2UDB_V81.

The EJB deployment tool does not run during installation of EJB 3.0 or later modules.

EJB references

An enterprise bean (EJB) reference is a logical name used to locate the home interface of an enterprise bean. EJB references are specified during deployment. At run time, EJB references are bound to the physical location (global JNDI name) of the enterprise beans in the target operational environment. EJB references are made available in the `java:comp/env/ejb` Java naming subcontext, or in another `java: namespace` if the reference name is a `java:module`, `java:app`, or `java:app URL`. EJB references with URL names are bound into the corresponding namespace according to the URL.

The product assigns default JNDI values for or automatically resolves incomplete EJB 3.0 reference targets.

For each EJB 2.1 or earlier EJB reference, you must specify a JNDI name. An example JNDI name for a *Supplier* EJB reference in a *Store* application might be `store/ejb/Supplier`. The binding definition is stored in IBM binding files such as `ibm-ejb-jar-bnd.xml`. When the referenced EJB is also deployed in the same application server, you can specify a server-scoped JNDI name. But if the referenced EJB is deployed on a different application server or if `ejb-ref` is defined in an application client module, then you should specify the global cell-scoped JNDI name.

If a deployer chooses to generate default bindings when installing the application, the installation wizard binds EJB references as follows: If an `<ejb-link>` is found, it is honored. If the `ejb-name` of an EJB defined in the application matches the `ejb-ref` name, then that EJB is chosen. Otherwise, if a unique EJB is found with a matching home (or local home) interface as the referenced bean, the reference is resolved automatically.

During and after application installation, you can specify EJB reference JNDI names on the Map EJB references to beans page. After installation, click **Applications > Application Types > WebSphere enterprise applications > application_name > EJB references** in the administrative console.

Note: To enable EJB reference targets to resolve automatically if the references are from EJB 2.1 or earlier modules or from Web 2.3 or earlier modules, select **Generate default bindings** on the Preparing for application installation page or select **Allow EJB reference targets to resolve automatically** on the Select installation options, Map EJB references to beans, or EJB references console pages.

Resource references

A resource reference is a logical name used to locate an external resource for an application. Resource references are specified during deployment. At run time, the references are bound to the physical location (global JNDI name) of the resource in the target operational environment. Resource references that do not use a URL for the JNDI name are made available as follows:

Table 28. Resource reference subcontexts. JNDI `java:comp/env` names are used for resource reference subcontexts.

Resource reference type	Subcontext declared in
Java DataBase Connectivity (JDBC) data source	<code>java:comp/env/jdbc</code>
JMS connection factory	<code>java:comp/env/jms</code>
JavaMail connection factory	<code>java:comp/env/mail</code>

Table 28. Resource reference subcontexts (continued). JNDI `java:comp/env` names are used for resource reference subcontexts.

Uniform Resource Locator (URL) connection factory	<code>java:comp/env/url</code>
---	--------------------------------

Applications alternatively can assign names to resource references that are `java:` URLs with prefixes such as `java:module`, `java:app`, and `java:global`. The URLs map to namespaces other than the component namespace, which contains `java:comp/env` name bindings. Resource references with URL names are bound into the corresponding namespace according to the URL.

For each resource reference, you must specify a JNDI name. If a deployer chooses to generate default bindings when installing the application, the installation wizard generates resource reference bindings derived from the `<res-ref-name>` tag, assuming that the `java:comp/env` name is the same as the resource global JNDI name.

During application installation, you can specify resource reference JNDI names on the Map resource references to references page. Specify JNDI names for the resources that represent the logical names defined in resource references. You can optionally specify login configuration name and authentication properties for the resource. After specifying authentication properties, click **OK** to save the values and return to the mapping step. Each resource reference defined in an application must be bound to a resource defined in your WebSphere Application Server configuration. After installation, click **Applications > Application Types > WebSphere enterprise applications > application_name > Resource references** in the administrative console to access the Resource references page.

Virtual host bindings for web modules

You must bind each web module to a specific virtual host. The binding informs a web server plug-in that all requests that match the virtual host must be handled by the web application. An example virtual host to be bound to a *Store* web application might be `store_host`. The binding definition is stored in IBM binding files such as `WEB-INF/ibm-web-bnd.xmi`.

If a deployer chooses to generate default bindings when installing the application, the installation wizard sets the virtual host to `default_host` for each `.war` file.

During and after application installation, you can map a virtual host to a web module defined in your application. On the Map virtual hosts for web modules page, specify a virtual host. The port number specified in the virtual host definition is used in the URL that is used to access artifacts such as servlets and JavaServer Pages (JSP) files in the web module. For example, an external URL for a web artifact such as a JSP file is `http://host_name:virtual_host_port/context_root/jsp_path`. After installation, click **Applications > Application Types > WebSphere enterprise applications > application_name > Virtual hosts** in the administrative console.

Message-driven beans

For each message-driven bean, you must specify a queue or topic to which the bean will listen. A message-driven bean is invoked by a Java Messaging Service (JMS) listener when a message arrives on the input queue that the listener is monitoring. A deployer specifies a listener port or JNDI name of an activation specification as defined in a connector module (`.rar` file) under **WebSphere Bindings** on the Beans page of an assembly tool EJB deployment descriptor editor. An example JNDI name for a listener port to be used by a *Store* application might be `StoreMdbListener`. The binding definition is stored in IBM bindings files such as `ibm-ejb-jar-bnd.xmi`.

If a deployer chooses to generate default bindings when installing the application, the installation wizard assigns JNDI names to incomplete bindings.

- For EJB 2.0 or later message-driven beans deployed as JCA 1.5-compliant resources, the installation wizard assigns JNDI names corresponding to `activationSpec` instances in the form `eis/MDB_ejb-name`.
- For EJB 2.0 or later message-driven beans deployed against listener ports, the listener ports are derived from the message-driven bean `<ejb-name>` tag with the string `Port` appended.

During application installation using the administrative console, you can specify a listener port name or an activation specification JNDI name for every message-driven bean on the Bind listeners for message-driven beans page. A listener port name must be provided when using the JMS providers: Version 5 default messaging, WebSphere MQ, or generic. An activation specification must be provided when the application's resources are configured using the default messaging provider or any generic J2C resource adapter that supports inbound messaging. If neither is specified, then a validation error is displayed after you click **Finish** on the Summary page.

After application installation, you can specify JNDI names and configure message-driven beans on console pages under **Resources > JMS > JMS providers** or under **Resources > Resource adapters**.

Restriction: You can only bind message driven-beans that are defined in an EJB 3.0 or later module to an activation specification.

Message destination references

A message destination reference is a logical name used to locate an enterprise bean in an EJB module that acts as a message destination. Message destination references exist only in J2EE 1.4 and later artifacts such as--

- J2EE 1.4 application clients
- EJB 2.1 projects
- 2.4 web applications

If multiple message destination references are associated with a single message destination link, then a single JNDI name for an enterprise bean that maps to the message destination link, and in turn to all of the linked message destination references, is collected during deployment. At run time, the message destination references are bound to the administered message destinations in the target operational environment.

If a message destination reference and a message-driven bean are linked by the same message destination, both the reference and the bean should have the same destination JNDI name. When both have the same name, only the destination JNDI name for the message-driven bean is collected and applied to the corresponding message destination reference.

If a deployer chooses to generate default bindings when installing the application, the installation wizard assigns JNDI names to incomplete message destination references as follows: If a message destination reference has a `<message-destination-link>`, then the JNDI name is set to `ejs/message-destination-linkName`. Otherwise, the JNDI name is set to `eis/message-destination-refName`.

Depending on the references in and artifacts used by your application, you might need to define bindings for references and artifacts not listed in this topic.

Application resource conflicts

Before Version 8 of the product, application-defined resources such as references and environment entries were bound into the component namespace relative to `java:comp/env`. In Version 8, an application developer can assign a name to a resource that is a `java:URL` prefixed with `java:module`, `java:app`, or `java:global`. Each of these URLs resolves to distinct namespaces, different from the component namespace. A `java:module` namespace is shared among all components in a module, a `java:app` namespace is shared among all modules in an application, and a `java:global` namespace is shared among all applications in a cell. Because the namespaces are shared, different resources might be assigned the same name, resulting in conflicts.

Conflicts at the module scope can occur only if two components in the module define resources with the same name. Because of the small size of this scope, it is unlikely for a module to have true conflicts. However, if multiple instances of the same resource definition exist, they must be configured the same. For

example, two EJB references to a particular EJB type which are both assigned the same `java:module` name must both be configured with the same binding data. Otherwise, the two resources will conflict and the application configuration action will fail.

Application-scoped resources are like module-scoped resources. The only difference is that the definitions can originate from any module in the application. As with module-scoped resources, all application-scoped resources that have the same name must be the same type of resource and must be configured with the same binding data.

Global-scoped resources differ from application- and module-scoped resources in that conflicts can occur among different applications. When a conflict occurs, conflicting applications cannot coexist if the resources are not logically the same. If multiple occurrences of a global-scoped resource all identify logically the same resource, they must all be configured with the same binding data in order to not be detected by the product as conflicting. To edit a global-scoped resource for which occurrences for multiple applications exist, all defining applications must be edited in the same session so as not to introduce a conflict. Failure to do so will result in a failure when the session is saved.

Enterprise application collection

Use this page to view and manage enterprise applications.

This page lists installed enterprise applications. System applications, which are central to the product, are not shown in the list because users cannot edit them. Examples of system applications include *isclite*, *managementEJB* and *filetransfer*.

To view this administrative console page, click **Applications > Application Types > WebSphere enterprise applications**.

To view the values specified for an application's configuration, click the application name in the list. The displayed application settings page shows the values specified. On the settings page, you can change existing configuration values and link to additional console pages that assist you in configuring the application.

To manage an installed enterprise application, enable the **Select** check box beside the application name in the list and click a button:

Table 29. Button descriptions. Use the buttons to manage enterprise applications.

Button	Resulting action
Start	Attempts to run the application. After the application starts up successfully, the state of the application changes to <i>Started</i> if the application starts up on all deployment targets, else the state changes to <i>Partial Start</i> .
Stop	Attempts to stop the processing of the application. After the application stops successfully, the state of the application changes to <i>Stopped</i> if the application stops on all deployment targets, else the state changes to <i>Partial Stop</i> .
Install	Opens a wizard that helps you deploy an application or a module such as a <code>.jar</code> , <code>.war</code> , <code>.sar</code> or <code>.rar</code> file onto a server or a cluster.
Uninstall	Deletes the application from the product configuration repository and deletes the application binaries from the file system of all nodes where the application modules are installed after the configuration is saved and synchronized with the nodes.
Update	Opens a wizard that helps you update application files deployed on a server. You can update the full application, a single module, a single file, or part of the application. If a new file or module has the same relative path as a file or module already existing on the server, the new file or module replaces the existing file or module. If the new file or module does not exist on the server, it is added to the deployed application.

Table 29. Button descriptions (continued). Use the buttons to manage enterprise applications.

Button	Resulting action
Rollout Update	<p>Sequentially updates an application installed on multiple cluster members across a cluster. After you update an application's files or configuration, click Rollout Update to install the application's updated files or configuration on all cluster members of a cluster on which the application is installed. Rollout Update does the following for each cluster member in sequence:</p> <ol style="list-style-type: none"> 1. Saves the updated application configuration. 2. Stops all of the cluster members on one node. 3. Updates the application on the node by synchronizing the configuration. 4. Restarts the stopped cluster members. 5. Repeats steps 2 through 4 for all of the nodes that have cluster members. <p>Use Rollout Update if the application is deployed on one or more clusters spread across multiple nodes. This action reduces the amount of time that any single cluster member is unavailable to serve requests to the smallest interval possible. Pending IIOP transactions will complete before a cluster member stops; in-flight HTTP and JMS transactions might be lost while the cluster member is stopping. For an application server without clusters, use Update and then save and synchronize the node instead. For a stand-alone application server, simply update and save.</p>
Remove File	<p>Deletes a file of the deployed application or module. Remove File deletes a file from the configuration repository and from the file system of all nodes where the file is installed.</p> <p>If the application or module is deployed on a cluster, after removing a file click Rollout Update to roll out the changes across the entire cluster.</p>
Export	<p>Accesses the Export Application EAR files page, which you use to export an enterprise application to an EAR file at a location of your choice. Use the Export action to back up a deployed application and to preserve its binding information.</p>
Export DDL	<p>Accesses the Export Application DDL files page, which you use to export DDL files (Table.ddl) in the EJB modules of an enterprise application to a location of your choice.</p>
Export File	<p>Accesses the Export a file from an application page, which you use to export a file of an enterprise application or module to a location of your choice.</p> <p>If the browser does not prompt for a location to store the file, click File > Save as and specify a location to save the file that is shown in the browser.</p>

These buttons are not available when you access this page from an application server settings page. When this page is accessed from an application server settings page, it is entitled the Installed applications page.

When security is enabled, a separate application list is shown for each of your administrative roles. Supported roles include monitor, configurator, operator, administrator, deployer, and administrative security manager. For example, when you have the administrator role, the statement "You can administer the following resources" is shown followed by a list of applications that you can administer.

Name

Specifies the name of the installed (or deployed) application. Application names must be unique within a cell and cannot contain an unsupported character.

Application Status

Indicates whether the application deployed on the application server is started, stopped, or unknown.

Table 30. Application status. The status indicates whether the application is running.







	Started	Application is running.
	Partial Start	Application is in the process of changing from a <i>Stopped</i> state to a <i>Started</i> state. Application is starting to run but is not fully running yet. Or, it cannot fully start because a server mapped to one or more application modules is stopped.

Table 30. Application status (continued). The status indicates whether the application is running.

	Stopped	Application is not running.
	Partial Stop	Application is in the process of changing from a <i>Started</i> state to a <i>Stopped</i> state. Application has not stopped running yet.
	Unknown	Status cannot be determined.
	Pending	Status is temporarily unknown pending an event that a user did not initiate, such as pending an asynchronous call.
	Not applicable	Application does not provide information as to whether it is running.

The status of an application on a web server is always **Unknown**.

Startup order

Specifies the order in which applications are started when the server starts. The application with the lowest startup order is started first.

This table column is available only when this page is accessed from an application server settings page; thus when this page is entitled the Installed applications page.

Enterprise application settings

Use this page to configure an enterprise application.

To view this administrative console page, click **Applications > Application Types > WebSphere enterprise applications > application_name**.

If you have a JAX-WS web service application installed, you also can click **Services > Service providers > service_name** or **Services > Service clients > service_name**.

Attention: If an application is running, changing an application setting causes the application to restart. On stand-alone servers, the application restarts after you save the change. On multiple-server products, the application restarts after you save the change and files synchronize on the node where the application is installed. To control when synchronization occurs on multiple-server products, deselect **Synchronize changes with nodes** on the Console preferences page.

Name:

Specifies a logical name for the application. An application name must be unique within a cell and cannot contain an unsupported character.

An application name cannot begin with a period (.), cannot contain leading or trailing spaces, and cannot contain any of the following characters:

Table 31. Characters that you cannot use in a name. The product does not support these characters in a name.

Unsupported characters		
/ forward slash	\$ dollar sign	' single quote mark
\ backslash	= equal sign	" double quote mark
* asterisk	% percent sign	vertical bar
, comma	+ plus sign	< left angle bracket
: colon	@ at sign	> right angle bracket
; semi-colon	# hash mark	& ampersand (and sign)
? question mark]]> No specific name exists for this character combination	

Data type String

Application reference validation:

Specifies whether the product examines the application references specified during application installation or updating and, if validation is enabled, warns you of incorrect references or fails the operation.

An application typically refers to resources using data sources for container managed persistence (CMP) beans or using resource references or resource environment references defined in deployment descriptors. The validation checks whether the resource referred to by the application is defined in the scope of the deployment target of that application.

The resource can be defined on the server, its node, cell or the cluster if the server belongs to a cluster. Select **Don't validate** for no resource validation, **Issue warnings** for warning messages about incorrect resource references, or **Stop installation if validation fails** to stop operations that fail as a result of incorrect resource references.

This **Application reference validation** setting is the same as the **Validate input off/warn/fail** field on the application installation and update wizards.

Data type String
Default Issue warnings

Configuring application startup

You can configure the startup behavior of an application. The values set affect how quickly an application starts and what occurs when an application starts.

Before you begin

This topic assumes that your application or module is already deployed on a server.

This topic also assumes that your application or module is configured to start automatically when the server starts. By default, an installed application starts when the server on which the application resides starts.

About this task

This topic describes how to change the settings of an application or module using the administrative console.

Attention: If an application is running, changing an application setting causes the application to restart. On stand-alone servers, the application restarts after you save the change. On multiple-server products, the application restarts after you save the change and files synchronize on the node where the application is installed. To control when synchronization occurs on multiple-server products, deselect **Synchronize changes with nodes** on the Console preferences page.

Procedure

1. Click **Applications > Application Types > WebSphere enterprise applications > *application_name* > Startup behavior** in the console navigation tree.
The Startup behavior settings page is displayed.
2. Specify the startup order for the application.
If your application starts automatically when its server starts, the value for **Startup order** on the Startup behavior settings page specifies the order in which applications are started when the server

starts. The application with the lowest startup order, or starting weight, is started first. For example, specify 1 for **Startup order** for applications that you want started first, specify 2 for applications to be started next, and so.

best-practices: For Session Initiation Protocol (SIP) applications, the `<load-on-startup>` tag in the `sip.xml` file affects the order in which applications are started. The value that you set for **Startup order** on the Startup behavior settings page determines the importance or weight of an application within a composition of SIP applications. For example, for the most important SIP application within a SIP application composition, specify 1 for **Startup order**. For the next most important SIP application within the composition, specify 2 for **Startup order**, and so on.

3. Specify whether the application must initialize fully before its server is considered started.
If your application starts automatically when its server starts, **Launch application before server completes startup** specifies whether the application must initialize fully before its server is considered started. Background applications can be initialized on an independent thread, thus allowing the server startup to complete without waiting for the application. This setting applies only if the application is run on a Version 6.0 or later application server.
4. Specify whether to create MBeans for resources such as servlets or JavaServer Pages (JSP) files within an application when the application starts.

The default for **Create MBeans for resources** is to create MBeans.

Results

The application or module configuration is changed. The application or stand-alone web module is restarted so the changes take effect.

What to do next

If the application or module is deployed on a cluster and you have no more configuration changes to make, click **Rollout Update** on the Enterprise applications page to propagate the changed configuration on all cluster members of the cluster on which the application or module is deployed. **Rollout Update** sequentially updates the configuration on the nodes that contain cluster members.

Save changes to your administrative configuration.

The application binaries are transferred to nodes when the configuration changes on the deployment manager synchronize with configurations for individual nodes on which the application will run.

Startup behavior settings

Use this page to configure when an application starts compared to other applications and to the server, and to configure whether MBeans for resources are created when an application starts.

To view this administrative console page, click **Applications > Application Types > WebSphere enterprise applications > *application_name* > Startup behavior**.

Attention: If an application is running, changing an application setting causes the application to restart. On stand-alone servers, the application restarts after you save the change. On multiple-server products, the application restarts after you save the change and files synchronize on the node where the application is installed. To control when synchronization occurs on multiple-server products, deselect **Synchronize changes with nodes** on the Console preferences page.

Startup order:

Specifies the order in which applications are started when the server starts. The startup order is like a starting weight. The application with the lowest starting weight is started first.

Note: Be aware of the following limitations:

- The first application might not start completely before the application server tries to start the next application in the startup order. This scenario might cause unexpected results if the second application depends on the first application.
- For Session Initiation Protocol (SIP) applications, the `<load-on-startup>` tag in the `sip.xml` file affects the order in which servlets within applications are started. The value that you set for **Startup order** on this Startup behavior console page determines the importance or weight of an application within a composition of SIP applications. For example, for the most important SIP application within a SIP application composition, specify 1 for **Startup order**. For the next most important SIP application within the composition, specify 2 for **Startup order**, and so on. For more information, see the JSR 116 specification.

Data type	Integer
Default	1
Range	0 to 2147483647

Launch application before server completes startup:

Specifies whether the application must initialize fully before the server starts.

The default setting of `false` indicates that server startup will not complete until the application starts.

A setting of `true` informs the product that the application might start on a background thread and thus server startup might continue without waiting for the application to start. Thus, the application might not be ready for use when the application server starts.

gotcha: Specifying this startup behavior setting is not valid for an application that is mapped to an application server that runs on the z/OS platform.

Data type	Boolean
Default	<code>false</code>

Create MBeans for resources:

Specifies whether to create MBeans for various resources, such as servlets or JavaServer Pages (JSP) files, within an application when the application starts. The default is to create MBeans.

Data type	Boolean
Default	<code>true</code>

Configuring binary location and use

You can designate where binary files (binaries) used by your application reside, whether the product distributes binaries for you automatically, and otherwise configure the use of binaries.

Before you begin

This topic assumes that your application or module is already deployed on a server.

About this task

This topic describes how to change the settings of an application or module using the administrative console.

Attention: If an application is running, changing an application setting causes the application to restart. On stand-alone servers, the application restarts after you save the change. On multiple-server products, the application restarts after you save the change and files synchronize on the node where the application is installed. To control when synchronization occurs on multiple-server products, deselect **Synchronize changes with nodes** on the Console preferences page.

Procedure

1. Click **Applications > Application Types > WebSphere enterprise applications > application_name > Application binaries** in the console navigation tree.

The Application binaries page is displayed.

2. Specify the directory to hold the application binaries.

The default is `${APP_INSTALL_ROOT}/cell_name`, where the `${APP_INSTALL_ROOT}` variable is `profile_root/installedApps`. For example:

```
C:/WebSphere/AppServer/profiles/profile_name/installedApps/cell_name
```

Refer to “Application binary settings” on page 297 for a detailed description of the **Location (full path)** setting.

3. Specify the bindings, extensions, and deployment descriptors that an application server uses.

By default, an application server uses the bindings, extensions, and deployment descriptors located with the application deployment document, the `deployment.xml` file.

To specify that the application server use the bindings, extensions, and deployment descriptors located in the application archive (EAR) file, select **Use configuration information in binary**. Select this setting for applications installed on 6.x or later deployment targets.

4. Specify whether the product distributes application binaries automatically to other nodes on the cell.

By default, **Enable binary distribution, expansion and cleanup post uninstallation** is selected and binaries are distributed automatically.

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.

Important: If you disable this option and you do not copy and expand the application binaries to the nodes, a later saving of the configuration or manual synchronization does not move the application binaries to the nodes for you.

5. Specify access permissions for binaries.

- a. Ensure that the **Enable binary distribution, expansion and cleanup post uninstallation** option is enabled. That option must be enabled to specify access permissions for binaries.

- b. For **File permissions**, specify a string that defines access permissions for binaries that are expanded in the named location.

You can specify file permissions in the text field. You can also set some of the commonly used file permissions by selecting them from the multiple-selection list. List selections overwrite file permissions set in the text field.

For details on **File permissions**, refer to “Application binary settings” on page 297.

6. Click **OK**.

Results

The application or module configuration is changed. The application or stand-alone web module is restarted so the changes take effect.

What to do next

If the application or module is deployed on a cluster and you have no more configuration changes to make, click **Rollout Update** on the Enterprise applications page to propagate the changed configuration

on all cluster members of the cluster on which the application or module is deployed. **Rollout Update** sequentially updates the configuration on the nodes that contain cluster members.

Save changes to your administrative configuration.

In multiple-server products, the application binaries are transferred to nodes when the configuration changes on the deployment manager synchronize with configurations for individual nodes on which the application will run.

Application binary settings

Use this page to configure the location and distribution of application binary files.

To view this administrative console page, click **Applications > Application Types > WebSphere enterprise applications > *application_name* > Application binaries**.

Attention: If an application is running, changing an application setting causes the application to restart. On stand-alone servers, the application restarts after you save the change. On multiple-server products, the application restarts after you save the change and files synchronize on the node where the application is installed. To control when synchronization occurs on multiple-server products, deselect **Synchronize changes with nodes** on the Console preferences page.

Location (full path):

Specifies the directory to which the enterprise application archive (EAR) file is installed. This **Location** setting is the same as the **Directory to install application** field on the application installation and update wizards.

By default, an EAR file is installed in the *profile_root/installedApps/cell_name/application_name.ear* directory.

Setting options include the following:

- Do not specify a value and leave the field empty.

The default value is *\${APP_INSTALL_ROOT}/cell_name*, where the *\${APP_INSTALL_ROOT}* variable is *profile_root/installedApps*. A directory having the EAR file name of the application being installed is appended to *\${APP_INSTALL_ROOT}/cell_name*. Thus, if you do not specify a directory, the EAR file is installed in the *profile_root/installedApps/cell_name/application_name.ear* directory.

- Specify a directory.

If you specify a directory, the application is installed in *specified_path/application_name.ear* directory. A directory having the EAR file name of the application being installed is appended to the path that you specified for **Directory to install application** when installing the application. For example, if you installed *Clock.ear* and specify *C:/myapps* on Windows machines, the application is installed in the *myapps/Clock.ear* directory. The *\${APP_INSTALL_ROOT}* variable is set to the specified path.

- Specify *\${APP_INSTALL_ROOT}/\${CELL}* for the initial installation of the application.

If you intend to export the application from one cell and later install the exported application on a different cell, specify the *\${CELL}* variable for the initial installation of the application. For example, specify *\${APP_INSTALL_ROOT}/\${CELL}* for this setting. Exporting the application creates an enhanced EAR file that has the application and its deployment configuration. The deployment configuration retains the cell name of the initial installation in the destination directory unless you specify the *\${CELL}* variable. Specifying the *\${CELL}* variable ensures that the destination directory has the current cell name, and not the original cell name.

Important: If an installation directory is not specified when an application is installed on a single-server configuration, the application is installed in *\${APP_INSTALL_ROOT}/cell_name*. When the server is made a part of a multiple-server 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 `${APP_INSTALL_ROOT}/cell_name`. However, an application that is installed after the server is added to the network configuration uses the default installation directory `${APP_INSTALL_ROOT}/network_cell_name`. To move the application to the `${APP_INSTALL_ROOT}/network_cell_name` location upon running the `addNode` operation, explicitly specify the installation directory as `${APP_INSTALL_ROOT}/${CELL}` during installation. In such a case, the application files can always be found under `${APP_INSTALL_ROOT}/current_cell_name`.

- If the application has been exported and you want to install the exported EAR file in a different cell or location, specify `${APP_INSTALL_ROOT}/cell_name/application_name.ear` if you did not specify `${APP_INSTALL_ROOT}/${CELL}` for the initial installation.

The exported EAR file is an enhanced EAR file that has the application and its deployment configuration. The deployment configuration retains the value for **Directory to install application** that was used for the previous installation of the application. Unless you specify a different value, the enhanced EAR file will be installed to the same directory as for the previous installation.

If you did not specify the `{CELL}` variable during the initial installation, the deployment configuration uses the cell name of the initial installation in the destination directory. If you are installing on a different cell, specify `${APP_INSTALL_ROOT}/cell_name/application_name.ear`, where `cell_name` is the name of the cell to which you want to install the enhanced EAR file. If you do not designate the current cell name, `cell_name` will be the original cell name even though you are installing the enhanced EAR file on a cell that has a different name.

- Specify an absolute path or a use pathmap variable.

You can specify an absolute path or use a pathmap variable such as `{MY_APPS}`. You can use a pathmap variable in any installation.

A pathmap variable 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, which you can access by clicking **Environment > WebSphere variables** in the administrative console.

Data type	String
Units	Full path name

Use configuration information in binary:

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 EAR file.

The default (false) is to use the binding, extensions, and deployment descriptors located in `deployment.xml`. To use the binding, extensions, and deployment descriptors located in the EAR file, enable this setting (true).

This **Use configuration information in binary** setting is the same as the **Use binary configuration** field on the application installation and update wizards. Select this setting for applications installed on 6.x or later deployment targets only.

Data type	Boolean
Default	false

Enable binary distribution, expansion and cleanup post uninstallation:

Specifies whether the product expands application binaries in the installation location during installation and deletes application binaries during uninstallation. The default is to enable application distribution. Application binaries for installed applications are expanded to the directory specified.

On single-server installations, the binaries are deleted when you uninstall and save changes to the configuration.

On multiple-server installations, the binaries are deleted when you uninstall and save changes to the configuration and synchronize changes.

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.

Important: If you disable this option and you do not copy and expand the application binaries to the nodes, a later saving of the configuration or manual synchronization does not move the application binaries to the nodes for you.

This **Enable binary distribution, expansion and cleanup post uninstallation** setting is the same as the **Distribute application** field on the application installation and update wizards.

Data type	Boolean
Default	true

File permissions:

Specifies access permissions for application binaries for installed applications that are expanded to the directory specified.

The **Enable binary distribution, expansion and cleanup post uninstallation** option must be enabled to specify file permissions.

You can specify file permissions in the text field. You can also set some of the commonly used file permissions by selecting them from the multiple-selection list. List selections overwrite file permissions set in the text field.

You can set one or more of the following file permission strings in the list. Selecting multiple options combines the file permission strings.

Table 32. File permission string sets for list options. Select a list option or specify a file permission string in the text field.

Multiple-selection list option	File permission string set
Allow all files to be read but not written to	.*=755
Allow executables to execute	.*\.dll=755#.*\.so=755#.*\.a=755#.*\.sl=755
Allow HTML and image files to be read by everyone	.*\.htm=755#.*\.html=755#.*\.gif=755#.*\.jpg=755

Instead of using the multiple-selection list to specify file permissions, you can specify a file permission string in the text field. File permissions use a string that has the following format:

file_name_pattern=permission#file_name_pattern=permission

where *file_name_pattern* is a regular expression file name filter (for example, *.**.jsp* for all JSP files), *permission* provides the file access control lists (ACLs), and # is the separator between multiple entries of *file_name_pattern* and *permission*. If # is a character in a *file_name_pattern* string, use \# instead.

If multiple file name patterns and file permissions in the string match a uniform resource identifier (URI) within the application, then the product uses the most stringent applicable file permission for the file. For example, if the file permission string is `.*\.\.jsp=775#a.*\.\.jsp=754`, then the `abc.jsp` file has file permission 754.

best-practices: Using regular expressions for file matching pattern compares an entire string URI against the specified file permission pattern. You must provide more precise matching patterns using regular expressions as defined by Java programming API. For example, suppose the following directory and file URIs are processed during a file permission operation:

Table 33. Example URIs for file permission operations. Results are shown following this table.

1	/opt/WebSphere/profiles/AppSrv01/installedApps/MyCell/MyApp.ear/MyWarModule.war
2	/opt/WebSphere/profiles/AppSrv01/installedApps/MyCell/MyApp.ear/MyWarModule.war/MyJsp.jsp
3	/opt/WebSphere/profiles/AppSrv01/installedApps/MyCell/MyApp.ear/MyWarModule.war/META-INF/MANIFEST.MF
4	/opt/WebSphere/profiles/AppSrv01/installedApps/MyCell/MyApp.ear/MyWarModule.war/WEB-INF/classes/MyClass.class
5	/opt/WebSphere/profiles/AppSrv01/installedApps/MyCell/MyApp.ear/MyWarModule.war/mydir/MyClass2.class
6	/opt/WebSphere/profiles/AppSrv01/installedApps/MyCell/MyApp.ear/MyWarModule.war/META-INF

The file pattern matching results are:

- `MyWarModule.war` does not match any of the URIs
- `.*MyWarModule.war.*` matches all URIs
- `.*MyWarModule.war$` matches only URI 1
- `.*\.\.jsp=755` matches only URI 2
- `.*META-INF.*` matches URIs 3 and 6
- `.*MyWarModule.war/.*/.*\.\.class` matches URIs 4 and 5

If you specify a directory name pattern for **File permissions**, then the directory permission is set based on the value specified. Otherwise, the **File permissions** value set on the directory is the same as its parent. For example, suppose you have the following file and directory structure:

```
/opt/WebSphere/profiles/AppSrv01/installedApps/MyCell/MyApp.ear/MyWarModule.war/MyJsp.jsp
```

and you specify the following file pattern string:

```
.*MyApp.ear$=755#.*\.\.jsp=644
```

The file pattern matching results are:

- Directory `MyApp.ear` is set to 755
- Directory `MyWarModule.war` is set to 755
- Directory `MyWarModule.war` is set to 755

best-practices: Regardless of the operation system, always use a forward slash (/) as a file path separator in file patterns.

Access permissions specified here are at the application level. You can also specify access permissions for application binaries in the node level configuration. The node level file permissions specify the maximum (most lenient) permissions that can be given to application binaries. Access permissions specified here at application level can only be the same as or more restrictive than those specified at the node level.

This setting is the same as the **File permission** field on the application installation and update wizards.

Data type String

Application build level:

Specifies an uneditable string that identifies the build version of the application.

Data type String

Configuring the use of class loaders by an application

You can configure whether your application and web modules use their own class loaders to load classes or use different class loaders, as well as configure the reloading of classes when application files are updated. Class loaders enable an application to access repositories of available classes and resources.

Before you begin

This topic assumes that your application or module is already deployed on a server.

The following note applies to the xmi file references in this topic:

Note: For IBM extension and binding files, the .xmi or .xml file name extension is different depending on whether you are using a pre-Java EE 5 application or module or a Java EE 5 or later application or module. An IBM extension or binding file is named `ibm-*-ext.xmi` or `ibm-*-bnd.xmi` where `*` is the type of extension or binding file such as `app`, `application`, `ejb-jar`, or `web`. The following conditions apply:

- For an application or module that uses a Java EE version prior to version 5, the file extension must be `.xmi`.
- For an application or module that uses Java EE 5 or later, the file extension must be `.xml`. If `.xmi` files are included with the application or module, the product ignores the `.xmi` files.

However, a Java EE 5 or later module can exist within an application that includes pre-Java EE 5 files and uses the `.xmi` file name extension.

The `ibm-webservices-ext.xmi`, `ibm-webservices-bnd.xmi`, `ibm-webservicesclient-bnd.xmi`, `ibm-webservicesclient-ext.xmi`, and `ibm-portlet-ext.xmi` files continue to use the `.xmi` file extensions.

About this task

Class loaders affect whether your application and its modules find the resources that they need to run effectively. You can select whether your application and web modules use their own class loaders to load classes, or use a parent class loader.

An application class loader groups Enterprise JavaBeans (EJB) modules, shared libraries, resource adapter archives (RAR files), and dependency Java archive (JAR) files associated to an application. Dependency JAR files are JAR files that contain code which can be used by both enterprise beans and servlets.

An application class loader is the parent of a web application archive (WAR) class loader. By default, a web module has its own WAR class loader to load the contents of the web module. The WAR class-loader policy value of an application class loader determines whether the WAR class loader or the application class loader is used to load the contents of the Web module.

You can also select whether classes are reloaded when application files are updated. For EJB modules or any non-web modules, enabling class reloading causes the application server run time to stop and start

the application to reload application classes. For web modules such as servlets and JavaServer Pages (JSP) files, a web container reloads a web module only when the IBM extension reloadingEnabled in the `ibm-web-ext.xml` file is set to true.

To configure use of class loaders by your application and web modules, use the Class loading and update detection page of the administrative console.

Attention: If an application is running, changing an application setting causes the application to restart. On stand-alone servers, the application restarts after you save the change. On multiple-server products, the application restarts after you save the change and files synchronize on the node where the application is installed. To control when synchronization occurs on multiple-server products, deselect **Synchronize changes with nodes** on the Console preferences page.

Procedure

1. Click **Applications > Application Types > WebSphere enterprise applications > *application_name* > Class loading and update detection** to access the Class loading and update detection page.

2. Specify whether to reload application classes when the application or its files are updated.

By default, class reloading is not enabled. Select **Override class reloading settings for web and EJB modules** to choose to reload application classes. You might specify different values for EJB modules and for web modules such as servlets and JSP files.

3. Specify the number of seconds to scan the application's file system for updated files.

The value specified for **Polling interval for updated files** takes effect only if class reloading is enabled. The default is the value of the reloading interval attribute in the IBM extension (`META-INF/ibm-application-ext.xml`) file of the enterprise application (EAR file). You might specify different values for EJB modules and for web modules such as servlets and JSP files.

To enable reloading, specify an integer value that is greater than zero (for example, 1 to 2147483647).

To disable reloading, specify zero (0).

4. Specify the class loader order for the application.

The application class loader order specifies whether the class loader searches in the parent class loader or in the application class loader first to load a class. The default is to search in the parent class loader before searching in the application class loader to load a class.

Select either of the following values for **Class loader order**:

Option	Description
Classes loaded with parent class loader first	Causes the class loader to search in the parent class loader first to load a class. This value is the standard for Development Kit class loaders and WebSphere Application Server class loaders.
Classes loaded with local class loader first (parent last)	Causes the class loader to search in the application class loader first to load a class. By specifying Classes loaded with local class loader first (parent last) , your application can override classes contained in the parent class loader. Attention: Specifying the Classes loaded with local class loader first (parent last) value might result in <code>LinkageErrors</code> or <code>ClassCastException</code> messages if you have mixed use of overridden classes and non-overridden classes.

5. Specify whether to use a single or multiple class loaders to load web application archives (WAR files) of your application.

By default, web modules have their own WAR class loader to load the contents of the `WEB-INF/classes` and `WEB-INF/lib` directories. The default WAR class loader value is `Class loader for each WAR file`

in application, which uses a separate class loader to load each WAR file. Setting the value to `Single class loader for application` causes the application class loader to load the web module contents as well as the EJB modules, shared libraries, RAR files, and dependency JAR files associated to the application. The application class loader is the parent of the WAR class loader.

Select either of the following values for **WAR class loader policy**:

Option	Description
Class loader for each WAR file in application	Uses a different class loader for each WAR file.
Single class loader for application	Uses a single class loader to load all of the WAR files in your application.

6. Click **OK**.

Results

The application or module configuration is changed. The application or stand-alone web module is restarted so the changes take effect.

What to do next

If the application or module is deployed on a cluster and you have no more configuration changes to make, click **Rollout Update** on the Enterprise applications page to propagate the changed configuration on all cluster members of the cluster on which the application or module is deployed. **Rollout Update** sequentially updates the configuration on the nodes that contain cluster members.

Save changes to your administrative configuration.

On multiple-server products, the application binaries are transferred to nodes when the configuration changes on the deployment manager synchronize with configurations for individual nodes on which the application will run.

Class loading and update detection settings

Use this page to configure use of class loaders by an application.

To view this administrative console page, click **Applications > Application Types > WebSphere enterprise applications > *application_name* > Class loading and update detection**.

Attention: If an application is running, changing an application setting causes the application to restart. On stand-alone servers, the application restarts after you save the change. On multiple-server products, the application restarts after you save the change and files synchronize on the node where the application is installed. To control when synchronization occurs on multiple-server products, deselect **Synchronize changes with nodes** on the Console preferences page.

Note: For IBM extension and binding files, the `.xmi` or `.xml` file name extension is different depending on whether you are using a pre-Java EE 5 application or module or a Java EE 5 or later application or module. An IBM extension or binding file is named `ibm-*-ext.xmi` or `ibm-*-bnd.xmi` where `*` is the type of extension or binding file such as `app`, `application`, `ejb-jar`, or `web`. The following conditions apply:

- For an application or module that uses a Java EE version prior to version 5, the file extension must be `.xmi`.
- For an application or module that uses Java EE 5 or later, the file extension must be `.xml`. If `.xmi` files are included with the application or module, the product ignores the `.xmi` files.

However, a Java EE 5 or later module can exist within an application that includes pre-Java EE 5 files and uses the `.xmi` file name extension.

The `ibm-webservices-ext.xml`, `ibm-webservices-bnd.xml`, `ibm-webservicesclient-bnd.xml`, `ibm-webservicesclient-ext.xml`, and `ibm-portlet-ext.xml` files continue to use the `.xml` file extensions.

Override class reloading settings for web and EJB modules:

Specifies whether to enable class reloading when application files are updated.

Select **Override class reloading settings for web and EJB modules** to set `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.

Reloading settings in the `deployment.xml` file override the reloading settings for all web and EJB modules that can be defined in `ibm-web-ext.xml` and `META-INF/ibm-application-ext.xml` files.

For JavaServer Pages (JSP) files in a web module, a web container reloads JSP files only when the IBM extension `jspReloadingEnabled` in the `jspAttributes` of the `ibm-web-ext.xml` file is set to `true`. You can enable JSP reloading during deployment on the JSP Reload Options page.

Data type	Boolean
Default	false

Polling interval for updated files:

Specifies the number of seconds to scan the application's file system for updated files. The default is the value of the reloading interval attribute in the IBM extension (`META-INF/ibm-application-ext.xml`) file of the EAR file.

This **Polling interval for updated files** setting is the same as the **Reload interval in seconds** field on the application installation and update wizards.

To enable reloading, specify a value greater than zero (for example, 1 to 2147483647). To disable reloading, specify zero (0). The range is from 0 to 2147483647.

The reloading interval attribute takes effect only if class reloading is enabled.

Data type	Integer
Units	Seconds
Default	3

Class loader order:

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 development kit class loaders and WebSphere Application Server class loaders is `Classes loaded with parent class loader first`. By specifying `Classes loaded with local class loader first (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 `Classes loaded with parent class loader first` and `Classes loaded with local class loader first (parent last)`. The default is to search in the parent class loader before searching in the application class loader to load a class.

For your application to use the default configuration of Jakarta Commons Logging in WebSphere Application Server, set this application class loader mode to `Classes loaded with parent class loader`

first. For your application to override the default configuration of Jakarta Commons Logging in WebSphere Application Server, your application must provide the configuration in a form supported by Jakarta Commons Logging and this class loader mode must be set to `Classes loaded with local class loader first` (parent last). Also, to override the default configuration, set the class loader mode for each web module in your application so that the correct logger factory loads.

Data type	String
Default	Classes loaded with parent class loader first

WAR class loader policy:

Specifies whether to use a single class loader to load all WAR files of the application or to use a different class loader for each WAR file.

The options are `Class loader for each WAR file in application` and `Single class loader for application`. The default is to use a separate class loader to load each WAR file.

Data type	String
Default	Class loader for each WAR file in application

Manage modules settings

Use this page to specify deployment targets where you want to install the modules that are contained in your application. Modules can be installed on the same deployment target or dispersed among several deployment targets.

On single-server products, a deployment target can be an application server or web server.

On multiple-server products, a deployment target can be an application server, cluster of application servers, or Web server.

To view this administrative console page, click **Applications > Application Types > WebSphere enterprise applications > *application_name* > Manage modules**. This page is similar to the Map modules to servers page on the application installation and update wizards.

On this page, each **Module** must map to one or more targets, identified under **Server**. To change a mapping:

1. In the list of mappings, select each module that you want mapped to the same target or targets.
2. If your security configuration supports multiple domains, select the domain that has the desired clusters or servers or keep the **All domains** default.
3. From the **Clusters and servers** list, select one or more targets. Select only appropriate deployment targets for a module. You cannot install modules that use WebSphere Application Server Version 8.x features on a Version 7.x or 6.x target server. Similarly, you cannot install modules that use Version 7.x features on a Version 6.x target server.

Use the Ctrl key to select multiple targets. For example, to have a web server serve your application, press the Ctrl key and then select an application server and the web server together. The product generates the plug-in configuration file, `plugin-cfg.xml`, for that web server based on the applications which are routed through it.

4. Click **Apply**.

Attention: If an application is running, changing an application setting causes the application to restart. On stand-alone servers, the application restarts after you save the change. On multiple-server products, the application restarts after you save the change and files synchronize on the node where the application is installed. To control when synchronization occurs on multiple-server products, deselect **Synchronize changes with nodes** on the Console preferences page.

If you accessed this Manage modules page from a console enterprise application page for an already installed application, you can also use this page to view and manage modules in your application.

To view the values specified for a module configuration, click the module name in the list. The displayed module settings page shows the values specified. On the settings page, you can change existing configuration values and link to additional console pages that assist you in configuring the module.

To manage a module, select the module name in the list and click a button:

Button	Resulting action
Remove	<p>Removes the selected module from the deployed application. The module is deleted from the application in the configuration repository and also from all of the nodes where the application is installed and running or expected to run.</p> <p>On multiple-server products, if the application is running on a node when the module file is deleted from the node as a result of configuration synchronization then the product stops the application, deletes the module file from the file system of the node, and restarts the application.</p>
Update	<p>Opens a wizard that helps you update modules in an application. If a module has the same URI as a module already existing in the application, the new module replaces the existing module. If the new module does not exist in the application, it is added to the deployed application.</p> <p>On multiple-server products, if the application is running on a node when the module file is updated on the node as a result of configuration synchronization then the product stops the application, updates the module file on the file system of the node, and restarts the application. If the application is running on a node when the module file is added as a result of configuration synchronization then the product starts the newly added module without stopping and restarting the running application.</p>
Remove File	<p>Deletes a file from a module of a deployed application.</p> <p>On multiple-server products, the file is also deleted from all the nodes where the module is installed after configuration is synchronized with nodes. If the application is running on a node when the module file is deleted from the node as a result of configuration synchronization then the product stops the application, deletes the module file from the file system of the node, and restarts the application.</p>
Export File	<p>Accesses the Export a file from an application page, which you use to export a file of an enterprise application or module to a location of your choice.</p> <p>If the browser does not prompt for a location to store the file, click File > Save as and specify a location to save the file that is shown in the browser.</p>

Display clusters and servers in the following domain

Lists the domains that your security configuration supports. This setting displays only on multiple-server products with a security configuration that supports multiple domains.

From this list, select the domain that has the clusters or servers on which you want to deploy an application or module. To see all available deployment targets, select **All Domains**.

Selecting a domain causes the **Clusters and servers** list to show only the deployment targets that are configured in the domain. You cannot deploy the modules in an application across deployment targets that belong to different security domains.

Clusters and servers

Lists the names of available deployment targets. This list is the same for every application that is installed in the cell.

From this list, select only appropriate deployment targets for a module. You must install an application, enterprise bean (EJB) module, Session Initiation Protocol (SIP) archive (SAR), web module, or client module on a Version 8.x target under any of the following conditions:

- The module supports Java Platform, Enterprise Edition (Java EE) 6.
- The module calls an 8.x runtime application programming interface (API).
- The module uses an 8.x product feature.

You must install an application, EJB, SAR, or web module on a Version 8.x or 7.x target under any of the following conditions:

- The module supports Java EE 5.
- The module calls a 7.x runtime API.
- The module uses a 7.x product feature.

If a module supports J2EE 1.4, then you must install the module on a Version 6.x, 7.x or 8.x deployment target. Modules that call a 6.1.x API or use a 6.1.x feature can be installed on a 6.1.x, 7.x or 8.x deployment target. Modules that require 6.1.x feature pack functionality can be installed on a 6.1.x deployment target that has been enabled with that feature pack or on a 7.x or 8.x deployment target.

You can install an application or module developed for a Version 5.x product on any deployment target.

Module

Specifies the name of a module in the installed (or deployed) application.

URI

Specifies the location of the module relative to the root of the application (EAR file).

Module type

Specifies the type of module, for example, a web module or EJB module.

This setting is shown on the Manage modules page accessed from a console enterprise application page.

Server

Specifies the name of each deployment target to which the module currently is mapped.

To change the deployment targets for a module, select one or more targets from the **Clusters and servers** list and click **Apply**. The new mapping replaces the previous mapping.

Mapping modules to servers

Each module of a deployed application must be mapped to one or more target servers. The target server can be an application server, cluster of application servers or web server.

Before you begin

You can map modules of an application or stand-alone Web module to one or more target servers during or after application installation using the console. This topic assumes that the module is already installed on a server and that you want to change the mappings.

Before you change a mapping, check the deployment targets. You must specify an appropriate deployment target for a module. Modules that use Version 8.x features cannot be installed onto Version 7.x or 6.x target servers. Similarly, modules that use Version 7.x features cannot be installed onto Version 6.x target servers.

About this task

During application installation, different deployment targets might have been specified.

You use the Manage modules page of the administrative console to view and change mappings. This page is displayed during application installation using the console and, after the application is installed, can be accessed from the enterprise application settings page.

On the Manage modules page, specify target servers where you want to install the modules contained in your application. Modules can be installed on the same application server or dispersed among several application servers. Also, specify the web servers as targets that will serve as routers for requests to your application. The plug-in configuration file, `plugin-cfg.xml`, for each web server is generated based on the applications which are routed through it.

Attention: If an application is running, changing an application setting causes the application to restart. On stand-alone servers, the application restarts after you save the change. On multiple-server products, the application restarts after you save the change and files synchronize on the node where the application is installed. To control when synchronization occurs on multiple-server products, deselect **Synchronize changes with nodes** on the Console preferences page.

Procedure

1. Click **Applications > Application Types > WebSphere enterprise applications > *application_name* > Manage modules** in the console navigation tree.
The Manage modules panel is displayed.
2. Examine the list of mappings.
Ensure that each **Module** entry is mapped to one or more targets, identified under **Server**.
3. Change a mapping as needed.
 - a. Select each module that you want mapped to the same targets.
In the list of mappings, select check boxes for the modules.
 - b. If your security configuration supports multiple domains, select the domain that has the desired clusters or servers, or keep the **All domains** default.
Selecting a domain causes the **Clusters and servers** list to show only the deployment targets that are configured in the domain. You cannot deploy the modules in an application across deployment targets that belong to different security domains.
 - c. From the **Clusters and servers** list, select one or more targets.
Select only appropriate deployment targets for a module. You cannot install modules that use WebSphere Application Server Version 8.x features on a Version 7.x or 6.x target server.
Use the Ctrl key to select multiple targets. For example, to have a web server serve your application, use the Ctrl key to select an application server or cluster and the Web server together to have the `plugin-cfg.xml` plug-in configuration file for that web server generated based on the applications that are routed through it.
 - d. Click **Apply**.
4. Repeat steps 2 and 3 until each module maps to the desired targets.
5. Click **OK**.

Results

The application or module configurations are changed. The application or stand-alone web module is restarted so the changes take effect.

Example

To install an application that has modules which support Java Platform, Enterprise Edition (Java EE) 5 or 6 to two servers, do the following:

1. Click the **Select All** icon to select all of the modules in the application.
2. While pressing Ctrl, select two Version 8 application servers from the **Clusters and servers** list.
3. Click **Apply**.
4. Click **OK**.

What to do next

If the application or module is deployed on a cluster and your configuration changes are complete, click **Rollout Update** on the Enterprise applications page to propagate the changed configuration on all cluster members of the cluster on which the application or module is deployed. **Rollout Update** sequentially updates the configuration on the nodes that contain cluster members.

Save changes to your administrative configuration.

The application binaries are transferred to nodes when the configuration changes on the deployment manager synchronize with configurations for individual nodes on which the application will run.

Note: If, at a future time, you intend to delete the server or cluster to which the modules are mapped, add a new server or cluster and then remap the modules to it before deleting the previous server or cluster. If you do not remap the modules to a new server or cluster before deleting the previous one, you must uninstall each application or module and then reinstall them on an existing server or cluster. After a server or cluster to which modules are mapped is deleted, the modules cannot be remapped to another existing server or cluster.

Mapping virtual hosts for web modules

A virtual host must be mapped to each web module of a deployed application. Web modules can be installed on the same virtual host or dispersed among several virtual hosts.

Before you begin

You can map a virtual host to a web module during or after application installation using the console. This topic assumes that the web module is already installed on a server and that you want to change the mappings.

Before you change a mapping, check the virtual hosts definitions. You can install a web module on any defined virtual host. To view information on previously defined virtual hosts, click **Environment > Virtual hosts** in the administrative console. Virtual hosts enable you to associate a unique port with a module or application. The aliases of a virtual host identify the port numbers defined for that virtual host. A port number specified in a virtual host alias is used in the URL that is used to access artifacts such as servlets and JavaServer Pages (JSP) files in a web module. For example, the alias `myhost:8080` is the `host_name:port_number` portion of the URL `http://myhost:8080/servlet/snoop`.

About this task

During application installation, a virtual host other than the one you want mapped to your web module might have been specified.

The default virtual host setting usually is `default_host`, which provides several port numbers through its aliases:

- 80** An internal, insecure port used when no port number is specified
- 9080** An internal port
- 9443** An external, secure port

Unless you want to isolate your web module from other modules or resources on the same node (physical machine), `default_host` is a suitable virtual host for your web module.

In addition to `default_host`, the product provides `admin_host`, which is the virtual host for the administrative console system application. `admin_host` is on port 9060. Its secure port is 9043. Do not select `admin_host` unless the web module relates to system administration.

Use the Virtual hosts page of the administrative console to view and change mappings. This page is displayed during enterprise application installation using the console and, after the application is installed, can be accessed from an enterprise application settings page.

On the Virtual hosts page, specify a virtual host for each web module. Web modules of an application can be installed on the same virtual host or on different virtual hosts.

Procedure

1. Click **Applications > Application Types > WebSphere enterprise applications > *application_name* > Virtual hosts** in the console navigation tree. The Virtual hosts page is displayed.
2. Examine the list of mappings. Ensure that each **Web module** entry has the desired virtual host mapped to it, identified under **Virtual host**.
3. Change the mappings as needed.
 - a. Select each web module that you want mapped to a particular virtual host. In the list of mappings, place a check mark in the **Select** check boxes beside the web modules.
 - b. From the **Virtual host** drop-down list, select the desired virtual host. If you selected more than one virtual host in step 1:
 - 1) Expand **Apply Multiple Mappings**.
 - 2) Select the desired virtual host from the **Virtual host** drop-down list.
 - 3) Click **Apply**.
4. Repeat steps 2 and 3 until a desired virtual host is mapped to each web module.
5. Click **OK**.

Results

The application or web module configurations are changed. The application or stand-alone web module is restarted so the changes take effect.

What to do next

After mapping virtual hosts, do the following:

1. If the application or module is deployed on a cluster and you have no more configuration changes to make, click **Rollout Update** on the Enterprise applications page to propagate the changed configuration on all cluster members of the cluster on which the application or module is deployed. **Rollout Update** sequentially updates the configuration on the nodes that contain cluster members.
2. Regenerate the plug-in configuration file.

- a. Click **Servers > Server Types > Web servers**.
 - b. Select the web server for which you want to generate a plug-in.
 - c. Click **Generate Plug-in**.
3. Save changes to your administrative configuration.

In multiple-server products, the application binaries are transferred to nodes when the configuration changes on the deployment manager synchronize with configurations for individual nodes on which the application will run.

Virtual hosts settings

Use this page to specify virtual hosts for web modules contained in your application. Web modules can be installed on the same virtual host or dispersed among several virtual hosts.

To view this administrative console page, click **Applications > Application Types > WebSphere enterprise applications > *application_name* > Virtual hosts**. This page is the same as the **Map virtual hosts for web modules** page on the application installation and update wizards.

On this page, each web module must map to a previously defined virtual host, identified under **Virtual host**. You can see information on previously defined virtual hosts by clicking **Environment > Virtual hosts** in the administrative console. Virtual hosts enable you to associate a unique port with a module or application. The aliases of a virtual host identify the port numbers defined for that virtual host. A port number specified in a virtual host alias is used in the URL that is used to access artifacts such as servlets and JavaServer Pages (JSP) files in a web module. For example, the alias `myhost:8080` is the `host_name:port_number` portion of the URL `http://myhost:8080/servlet/snoop`.

The default virtual host setting usually is `default_host`, which provides several port numbers through its aliases:

- 80** An internal, insecure port used when no port number is specified
- 9080** An internal port
- 9443** An external, secure port

Unless you want to isolate your web module from other modules or resources on the same node (physical machine), `default_host` is a suitable virtual host for your web module.

In addition to `default_host`, the product provides `admin_host`, which is the virtual host for the administrative console system application. `admin_host` is on port 9060. Its secure port is 9043. Do not select `admin_host` unless the web module relates to system administration.

To change a mapping:

1. In the list of mappings, select the **Select** check box beside each web module that you want mapped to a particular virtual host.
2. From the **Virtual host** drop-down list, select the desired virtual host. If you selected more than one virtual host in step 1:
 - a. Expand **Apply Multiple Mappings**.
 - b. Select the desired virtual host from the **Virtual Host** drop-down list.
 - c. Click **Apply**.
3. Click **OK**.

Attention: If an application is running, changing an application setting causes the application to restart. On stand-alone servers, the application restarts after you save the change. On multiple-server products, the application restarts after you save the change and files synchronize on the node where the application is installed. To control when synchronization occurs on multiple-server products, deselect **Synchronize changes with nodes** on the Console preferences page.

Web module:

Specifies the name of a web module in the application that you are installing or that you are viewing after installation.

Virtual host:

Specifies the name of the virtual host to which the Web module is currently mapped.

Expanding the drop-down list displays a list of previously defined virtual hosts. To change a mapping, select a different virtual host from the list.

Do not specify the same virtual host for different web modules that have the same context root and are deployed on targets belonging to the same node even if the web modules are contained in different applications. Specifying the same virtual host causes a validation error.

Mapping properties for a custom login or trusted connection configuration

Use this page to view and manage the mapping properties for a custom login configuration or a trusted connection configuration.

To access the administrative console panel, complete the following steps:

1. Click **Applications > Application types> WebSphere enterprise applications > *application_name***.
2. From Enterprise JavaBeans Properties, click **Map data sources for all 2.x CMP beans**.
3. For container authorization, modify the authorization type by selecting your Enterprise JavaBeans(EJB) module and selecting **Container** from the Resource authorization menu.
4. Click **Apply**.
5. From Specify authentication method, select **Use custom login configuration** or **Use trusted connections** and the name of the application login configuration.
6. Select the name of your EJB module.
7. Click **Apply**.
8. Click **Mapping properties** in the Resource authorization column. This property is not available until after you click **Apply** in the previous step.

Name

Specifies the name for the mapping property.

Do not use the MAPPING_ALIAS property name because the name is reserved by the product.

Value

Specifies the value paired with the specified name.

Description

Specifies additional information about the name and value pair.

Viewing deployment descriptors

A deployment descriptor is an extensible markup language (XML) file that specifies configuration and container options for an application or module.

Before you begin

This topic assumes that you have installed an application or module on a server and that you want to view its deployment descriptor.

About this task

When you create a Java 2 Platform, Enterprise Edition (J2EE) application or module in an assembly tool, the assembly tool creates deployment descriptor files for the application or module. Java Platform, Enterprise Edition (Java EE) 5 or later applications and modules might use annotations instead of deployment descriptors.

After an application or module is installed on a server, you can view its deployment descriptor in the administrative console. You cannot view Java EE 5 or later annotations.

Unless an application supports Java EE 5 or later, an enterprise archive (EAR) file must contain an `application.xml` file. The `application.xml` identifies each module of an application. A Java EE 5 application is not required to provide an `application.xml` file in the EAR file. When an `application.xml` file does not exist, the product examines the Java archive (JAR) file contents to determine whether the JAR file is an enterprise bean (EJB) module or an application client module. A JAR file should not contain more than one deployment descriptor in it. When an `ejb-jar.xml` file is found in a JAR file, the product considers it an EJB module. If an `ejb-jar.xml` file is not found and an `application-client.xml` is found, the product considers the JAR file to be an application client module. If both `ejb-jar.xml` and `application-client.xml` files exist in the JAR file, the product might consider a JAR file intended to be an application client module to be an EJB module or a JAR file intended to be an EJB module to be an application client module. A JAR file should not contain more than one kind of deployment descriptor.

Procedure

1. Access a deployment descriptor view.

Click the navigational option stated in **Accessing a console view** to view the deployment descriptor for a given module:

Table 34. Accessing View deployment descriptor pages in the console. Click links on the console navigation tree and pages.

Module	Deployment descriptor file	Accessing a console view
Enterprise application	<code>application.xml</code>	Applications > Application Types > WebSphere enterprise applications > <i>application_name</i> > View deployment descriptor
Web application	<code>WEB-INF/web.xml</code>	Applications > Application Types > WebSphere enterprise applications > <i>application_name</i> > Manage modules > <i>module_name</i> > View deployment descriptor
	<code>WEB-INF/portlet.xml</code>	Applications > Application Types > WebSphere enterprise applications > <i>application_name</i> > Manage modules > <i>module_name</i> > View portlet deployment descriptor
Enterprise bean	<code>ejb-jar.xml</code>	Applications > Application Types > WebSphere enterprise applications > <i>application_name</i> > Manage modules > <i>module_name</i> > View deployment descriptor
Application client	<code>application-client.xml</code>	Applications > Application Types > WebSphere enterprise applications > <i>application_name</i> > Manage modules > <i>module_name</i> > View deployment descriptor

Table 34. Accessing View deployment descriptor pages in the console (continued). Click links on the console navigation tree and pages.

Module	Deployment descriptor file	Accessing a console view
Web service	webservices.xml	<p>Applications > Application Types > WebSphere enterprise applications > <i>application_name</i> > Manage modules > <i>module_name</i> ></p> <ul style="list-style-type: none"> • View web services client deployment descriptor extension • View web services server deployment descriptor extension • View web services server deployment descriptor extension <p>For information about the views, see the topic on viewing Web services deployment descriptors in the administrative console.</p>
Resource adapter embedded in enterprise application	ra.xml	<p>Applications > Application Types > WebSphere enterprise applications > <i>application_name</i> > Manage modules > <i>ra_module_name</i> > View deployment descriptor</p>
Stand-alone Resource adapter	ra.xml	<p>Resources > Resource Adapters > Resource adapters > <i>module_name</i> > View deployment descriptor</p>

2. Click **Expand All** to view the deployment descriptor contents.

Results

The deployment descriptor for the application or module is displayed.

Example

The deployment descriptor for the product DefaultApplication follows:

```
<application id="Application_ID" >
  <display-name> DefaultApplication.ear</display-name>
  <description> This is the IBM WebSphere Application Server Default Application.</description>
  <module id="WebModule_1" >
    <web>
      <web-uri> DefaultWebApplication.war</web-uri>
      <context-root> /</context-root>
    </web>
  </module>
  <module id="EjbModule_1" >
    <ejb> Increment.jar</ejb>
  </module>
  <security-role id="SecurityRole_1204342979281" >
    <description> All Authenticated users role.</description>
    <role-name> All Role</role-name>
  </security-role>
</application>
```

What to do next

After displaying a deployment descriptor on the console page, do the following:

1. Examine the deployment descriptor contents, including any configurations that it has for application bindings, security roles, references to other resources, or Java Naming and Directory Interface (JNDI) names.

For example, examine the JAR files of your Java EE 5 or later module to ensure that each JAR file does not contain more than one kind of deployment descriptor. If a JAR file contains more than one kind of deployment descriptor, proceed to the next step and remove the extraneous deployment descriptor. Thus, if both ejb-jar.xml and application-client.xml files exist in a JAR file, remove the deployment descriptor that your module does not need.

2. Change a deployment descriptor as needed.

You can edit a deployment descriptor file manually. However, it is preferable to edit a deployment descriptor using the console or in an assembly tool deployment descriptor editor to ensure that the deployment descriptor has valid properties and that its references contain appropriate values.

If your Java EE 5 or later module does not have a `metadata-complete` attribute or the `metadata-complete` attribute is set to `false`, you can instruct the product to write the entire module deployment descriptor, including deployment information from annotations, to XML format. On the Metadata for modules page, select **metadata-complete attribute**.

Note: If your Java EE 5 or later application uses annotations and a shared library, do not select **metadata-complete attribute**. When your application uses annotations and a shared library, setting the `metadata-complete` attribute to `true` causes the product to incorrectly represent an @EJB annotation in the deployment descriptor as `<ejb-ref>` rather than `<ejb-local-ref>`. For web modules, setting the `metadata-complete` attribute to `true` might cause `InjectionException` errors. If you must set the `metadata-complete` attribute to `true`, avoid errors by not using a shared library, by placing the shared library in either the `classes` or `lib` directory of the application server, or by fully specifying the metadata in the deployment descriptors.

Metadata for module settings

Use this page to instruct a Java Platform, Enterprise Edition (Java EE) enterprise bean (EJB) deployment descriptor, web module deployment descriptor, or JCA resource adapter archive (RAR) module to ignore annotations that specify deployment information.

To view this administrative console page, click **Applications > Application Types > WebSphere enterprise applications > *application_name* > Metadata for modules**. This page is the same as the Metadata for modules page on the application installation and update wizards.

If your application contains Java EE 5 or later modules, you can select to lock the deployment descriptor of one or more of the modules on the Metadata for modules page. If you select a **metadata-complete attribute** check box (set the `metadata-complete` attribute to `true`) and lock deployment descriptors, the product writes the complete module deployment descriptor, including deployment information from annotations, to XML format.

Annotations are a standard mechanism of adding metadata to Java classes. You can use metadata to simplify development and deployment of Java EE 5 or later artifacts. Prior to the introduction of Java language annotations, deployment descriptors were the standard mechanism used by Java EE components. These deployment descriptors were mapped to XML format, which facilitated their persistence. If you select to lock deployment descriptors, the product merges Java EE annotation-based metadata with the XML-based existing deployment descriptor metadata and persists the result.

When applications contain a large number of Java classes, the deployment processing time for the annotations can increase. To minimize the performance impact, you can use one of the following methods:

- Determine whether the module needs to use Java EE 5 or 6. If the module does not need to use Java EE 5 or 6, the annotations within the Java classes are not scanned.
- Use the “`metadata-complete attribute`” on page 276 in the module descriptor if the module uses Java EE 5 or later and it does not contain any annotations. This attribute disables the annotations processing for the module, but Java EE 5 or later modules might still be placed in the descriptor file. If you are migrating your application, but you are not adding annotations, consider using this attribute value.
- Restructure the application to place the utility Java archive (JAR) files into shared libraries if those JAR files do not have annotation information. Consider this method if you cannot set the “`metadata-complete attribute`” on page 276.
- Move the JAR files in the `WEB-INF/lib` directory to the root directory of the enterprise archive (EAR) file. Nested archives, such as a JAR file that is within a web application archive (WAR) that is within an EAR file, are very cumbersome to search through because of the multiple levels of compression.

Module

Specifies the name of a module in the installed (or deployed) application.

Data type String

URI

Specifies the location of the module relative to the root of the EAR file.

Data type String

metadata-complete attribute

Specifies whether to write the complete module deployment descriptor, including deployment information from annotations, to extensible markup language (XML) format.

By default, a **metadata-complete attribute** check box is not selected and the product does not write out annotation data to a module deployment descriptor.

If your modules do not have a metadata-complete attribute or the metadata-complete attribute is set to false, you can select a check box and instruct the product to write out annotation data to a module deployment descriptor.

Note: If your Java EE 5 or later application uses annotations and a shared library, do not select **metadata-complete attribute**. When your application uses annotations and a shared library, setting the metadata-complete attribute to true causes the product to incorrectly represent an @EJB annotation in the deployment descriptor as <ejb-ref> rather than <ejb-local-ref>. For web modules, setting the metadata-complete attribute to true might cause InjectionException errors. If you must select **metadata-complete attribute** (set the metadata-complete attribute to true), avoid errors by not using a shared library, by placing the shared library in either the classes or lib directory of the application server, or by fully specifying the metadata in the deployment descriptors.

After you select a check box, you cannot deselect (clear) the check box and the module is no longer shown in the list of modules on this page. If you select all the check boxes, the link to this page is no longer shown on the enterprise application settings page.

Data type Boolean
Default false (deselected)

Starting or stopping enterprise 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).

Before you begin

This topic assumes that the Java Platform, Enterprise Edition (Java EE) application is installed on a server. By default, the application starts automatically when the server starts.

About this task

You can start and stop applications manually using the following:

- Administrative console
- startApplication and stopApplication attributes of the AdminControl object with the wsadmin tool
- startApplication and stopApplication administrative jobs of the AdminTask.submitJob -jobType object with the wsadmin tool

- Java programs that use ApplicationManager or AppManagement MBeans

This topic describes how to use the administrative console to start or stop an application.

Note: This topic applies to applications that do not contain Java Application Programming Interface (API) for XML-Based Web Services (JAX-WS) or Service Component Architecture (SCA) service providers. To stop or start applications that contain JAX-WS or SCA service providers, use the Service providers page accessed by clicking **Services > Service providers**. To start a service provider application, select a service and click **Start Application**. To stop a service provider application, select a service and click **Stop Application**. Then, on the Stop application page, click **OK** to stop all modules in the application, including other services such as enterprise beans and servlets.

Procedure

1. Go to the Enterprise applications page. Click **Applications > Application Types > WebSphere enterprise applications** in the console navigation tree.
2. Select the check box for the application you want started or stopped.
3. Click a button:

Option	Description
Start	Runs the application and changes the state of the application to Started. The status is changed to <i>partially started</i> if not all servers on which the application is deployed are running.
Stop	Stops the processing of the application and changes the state of the application to <i>Stopped</i> .

To restart a running application, select the application you want to restart, click **Stop** and then click **Start**.

Results

The status of the application changes and a message stating that the application started or stopped displays at the top the page.

If an application server on which the application is deployed synchronizes configuration with the deployment manager during server startup, then the application might not start and a `DeploymentDescriptorLoadException` might be written to the server `SystemErr.log` file. Stop and restart the server, and then try starting the application again.

Note: This topic references one or more of the application server log files. Beginning in WebSphere Application Server Version 8.0 you can configure the server to use the High Performance Extensible Logging (HPEL) log and trace infrastructure instead of using `SystemOut.log`, `SystemErr.log`, `trace.log`, and `activity.log` files or native z/OS logging facilities. If you are using HPEL, you can access all of your log and trace information using the LogViewer command-line tool from your server profile bin directory. See the information about using HPEL to troubleshoot applications for more information on using HPEL.

What to do next

You can configure an application so it does not start automatically when the server on which it resides starts. You then start the application manually using options described in this topic.

If you want your application to start automatically when its server starts, you can adjust values that control how quickly the application or its server starts:

1. Click **Applications > Application Types > WebSphere enterprise applications > *application_name* > Startup behavior**.
2. Specify a different value for **Startup order**.
This setting specifies the order in which applications are started when the server starts. The default value is 1 in a range from 0 to 2147483647. The application with the lowest starting weight is started first.
3. Specify a different value for **Launch application before server completes startup**.
This setting specifies whether the application must initialize fully before its server starts. The default value of `false` prevents the server from starting completely until the application starts. To reduce the amount of time it takes to start the server, you can set the value to `true` and have the application start on a background thread, thus allowing server startup to continue without waiting for the application.
4. If the application or module is deployed on a cluster and you have no more configuration changes to make, click **Rollout Update** on the Enterprise applications page to propagate the changed configuration on all cluster members of the cluster on which the application or module is deployed. **Rollout Update** sequentially updates the configuration on the nodes that contain cluster members.
5. Save the changes to the application configuration.

Disabling automatic starting of applications

You can enable and disable the automatic starting of an application. By default, an installed application starts automatically when the server on which the application resides starts.

Before you begin

This topic assumes that the application is installed on an application server and that the application starts automatically when the server starts.

This topic also assumes that you mapped the installed application to a server or cluster and that you have an administrative role with an authority higher than monitor.

About this task

You might want an application to run only after you start it manually and not to run every time after the server starts. The target mapping for an application controls whether an application starts automatically when the server starts or requires you to start the application manually.

You must have an administrative role with an authority higher than monitor to change the automatic starting setting.

Procedure

1. Go to the Target specific application status page for your application.
Click **Applications > Application Types > WebSphere enterprise applications > *application_name* > Target specific application status**.
2. Select the target server or cluster on which the application resides.
3. Click **Disable Auto Start**.
4. Save changes to the administrative configuration.

Results

The application does not start when its server starts. You must start the application manually.

What to do next

To enable automatic starting of the application, do the following:

1. On the Target specific application status page for the application, select the target on which the application resides.
2. Click **Enable Auto Start**.
3. Save changes to the configuration.

Target specific application status

Use this page to view mappings of deployed applications or modules to servers.

On multiple-server products, the target can be a cluster as well as a server.

Also use this page to enable or disable the automatic starting of an application when the server on which the application resides starts.

To view this administrative console page, click **Applications > Application Types > WebSphere enterprise applications > *application_name* > Target specific application status**.

When security is enabled, a separate application list is shown for each of your administrative roles. Supported roles include monitor, configurator, operator, administrator, deployer, and administrative security manager. For example, when you have the administrator role, the statement “You can administer the following resources” is shown followed by a list of servers that you can administer.

Target

States the name of the target server to which the application or module maps. You specify the target on the Manage modules page accessed from the settings for an application.

Node

Specifies the node name if the target is a server.

Version

Specifies the version level of the target. The target can be a Version 8.x, 7.x or 6.x deployment target.

A *deployment target* is a server with all members on a WebSphere Application Server product. For example, an *8.x deployment target* is a server with all members on a WebSphere Application Server Version 8.0 or later product.

An application, enterprise bean (EJB) module, Session Initiation Protocol (SIP) archive (SAR), web module, or client module must be installed on a Version 8.x target under any of the following conditions:

- The module supports Java Platform, Enterprise Edition (Java EE) 6.
- The module calls an 8.x runtime application programming interface (API).
- The module uses an 8.x product feature.

An application, EJB, SAR, or web module must be installed on a Version 8.x or 7.x target under any of the following conditions:

- The module supports Java EE 5.
- The module calls a 7.x runtime API.
- The module uses a 7.x product feature.

If a module supports Java 2 Platform, Enterprise Edition (J2EE) 1.4, then you can install the module on a Version 6.x or later deployment target. Modules that call a 6.1.x API or use a 6.1.x feature can be installed on a 6.1.x, 7.x or 8.x deployment target. Modules that require 6.1.x feature pack functionality can be installed on a 7.x or 8.x deployment target or on a 6.1.x deployment target that has been enabled with that feature pack.

If JavaServer Pages (JSP) precompilation, EJB deployment (ejbdeploy), or Web Services deployment (wsdeploy) are enabled, then you can deploy applications to only those targets that have same product version as the deployment manager. If applications are targeted to servers that have an earlier version

than the deployment manager, then you cannot deploy to those targets. Thus, if JSP precompilation, `ejbdeploy`, or `wsdeploy` are enabled, then you must deploy the application on a 6.1.x, 7.x or 8.x target.

You can install an application or module developed for a Version 5.x product on any deployment target.

Auto Start

Specifies whether the application modules installed on the target server are started (or enabled) when the server starts. This setting specifies the initial state of application modules. A **Yes** value indicates that the corresponding modules are enabled and thus are accessible when the server starts. A **No** value indicates that the corresponding modules are not enabled and thus are not accessible when the server starts.







By default, Auto Start is enabled. Thus, by default an installed application starts automatically when the server on which the application resides starts.

If you have an administrative role with an authority higher than monitor, you can enable and disable the automatic starting of the application. To disable the automatic starting of the application, enable the **Select** check box beside the target server and click **Disable Auto Start**. When automatic starting is disabled, the application does not start when its server starts. To enable the automatic starting of the application, select the target and click **Enable Auto Start**.

Application Status

Indicates whether the application deployed on the application server is started, stopped, or unknown.

Table 35. Application status. Shows whether the application is running.

	Started	Application is running.
	Partial Start	Application is in the process of changing from a <i>Stopped</i> state to a <i>Started</i> state. Application is starting to run but is not fully running yet. The application might be in the Partial Start state because one of its application servers is not started.
	Stopped	Application is not running.
	Partial Stop	Application is in the process of changing from a <i>Started</i> state to a <i>Stopped</i> state. Application has not stopped running yet.
	Unknown	Status cannot be determined.
	Pending	Status is temporarily unknown pending an event that a user did not initiate, such as pending an asynchronous call.
	Not applicable	Application does not provide information as to whether it is running.

The status of an application on a web server is always **Unknown**.

Updating enterprise application files

You can update Java Platform, Enterprise Edition (Java EE) application files deployed on a server.

Before you begin

Update your Java EE application or modules and reassemble them using an assembly tool. Typical tasks include adding or editing assembly properties, adding or importing modules into an application, and adding enterprise beans, web components, and files.

Also, determine whether the updated files can be installed to your deployment targets. Version 8.0 supports Java EE 6 enterprise applications and modules.

If you are deploying Java EE 6 modules, ensure that the deployment target and its node support Version 8.0. You can deploy Java EE 6 modules to Version 8.0 servers or to clusters that contain Version 8.0

cluster members only. You cannot deploy Java EE 6 modules to servers on Version 7.x or 6.x nodes or to clusters that contain Version 7.x or 6.x cluster members.

The administrative console Server collection pages show the versions for deployment targets.

About this task

Updating consists of adding a new file or module to an installed application, or replacing or removing an installed application, file or module. After replacement of a full application, the old application is uninstalled. After replacement of a module, file or partial application, the old installed module, file or partial application is removed from the installed application.

Service Component Architecture (SCA) applications can contain Java EE modules as well as SCA composites. You can use the procedure described in this topic to update Java EE modules in SCA applications. However, you cannot use this procedure to update SCA composites or other SCA-specific files.

Procedure

1. Determine which method to use to update your application files. The product provides several ways to update modules.
2. Update the application files using
 - Administrative console
 - Drag and drop to a monitored directory
 - wsadmin scripts
 - Java application programming interfaces
 - WebSphere rapid deployment of Java EE applications

In some situations, you can update applications or modules without restarting the application server using hot deployment. Do not use hot deployment unless you are an experienced user and are updating applications in a development or test environment.

3. If needed, restart the application manually so the changes take effect. Start the deployed application files using
 - Administrative console
 - wsadmin startApplication
 - Java programs that use ApplicationManager or AppManagement MBeans

When you update an application while it is running, the product automatically stops the application or only its changed components, updates the application logic, and restarts the stopped application or its components.

If you update module metadata while an application is running, restarting the application might not be sufficient for the changes to take effect. For example, if you change descriptors in running Java EE 6 applications that use annotations, you must reinstall the application. If you change classes that introduce, remove, or alter class hierarchies within an application, and those changes impact annotated classes, you also must reinstall the application.

What to do next

Save the changes to your administrative configuration.

When saving the configuration, synchronize the configuration with the nodes where the application is expected to run.

Next, test the application. For example, point a web browser at the URL for a deployed application (typically `http://hostname:9060/web_module_name`, where *hostname* is your valid web server and 9060 is the default port number) and examine the performance of the application. If the application does not

perform as desired, edit the application configuration, then save and test it again.

Ways to update enterprise application files

You can update Java Platform, Enterprise Edition (Java EE) application files deployed on a server or cluster in several ways.

Table 36. Ways to update application files. You can update application files using the console, wsadmin, programming, or deployment tools

Option	Method	Comments	Starting after update
<p>Administrative console update wizard</p> <p>See “Updating enterprise applications with the console” on page 324.</p> <p>To remove a single file from a Java EE application or module, see the topic on removing enterprise files.</p>	<p>Briefly, do the following:</p> <ol style="list-style-type: none"> 1. Go to the Enterprise applications page. Click Applications > Application Types > WebSphere enterprise applications in the console navigation tree. 2. Select the application to update and click Update. 3. On the Preparing for application update page, identify the application, module or files to update and click Next. 4. Complete steps in the update wizard and click Finish. 	<p>On the Preparing for application update page:</p> <ul style="list-style-type: none"> • Use Full application to update an .ear file. • Use Single module to update a .war, .sar, enterprise bean .jar, or connector .rar file. • Use Single file to update a file other than an .ear, .war, .sar, EJB .jar, or .rar file. • Use Partial application to update or remove multiple files. 	<p>On the Enterprise applications page, select the updated application and click Start.</p>
<p>Monitored directory</p>	<p>If you installed an application or module by adding it to a monitored directory, you can update it by adding an application or module that has the same file name to the monitored directory.</p>	<p>Refer to the topic on installing enterprise application files by adding them to a monitored directory.</p>	<p>After application or module installation, the product starts it automatically.</p> <p>For installations on a deployment manager monitored directory, the product does not start an application or module automatically if the node agent or server is not running.</p>
<p>Application properties files</p>	<p>Create a properties file that specifies to update application files. Then, run the wsadmin applyConfigProperties command in any of the following ways:</p> <ul style="list-style-type: none"> • In a script • At a command prompt • By adding the properties file to a dragDropDeployableApps/ deploymentProperties monitored directory 	<p>Refer to the topic on installing enterprise application files by adding properties files to a monitored directory.</p>	<p>After application or module installation, the product starts it automatically.</p> <p>For installations on a deployment manager monitored directory, the product does not start an application or module automatically if the node agent or server is not running.</p>

Table 36. Ways to update application files (continued). You can update application files using the console, wsadmin, programming, or deployment tools

Option	Method	Comments	Starting after update
wsadmin scripts See the topic on updating installed applications using the wsadmin scripting tool.	Use the update command or the updateInteractive command in a script or at a command prompt. For more information on the update and updateInteractive commands, see the topic on commands for the AdminApp object.	The Getting started with wsadmin scripting topic provides an overview of wsadmin.	Start the application using the invoke command and the startApplication attribute. For more information about the invoke command, see the topic on commands for the AdminControl object.
Java application programming interfaces See the topic on using administrative programs (JMX).	Update deployed applications by completing the steps in the topic on managing applications through programming.	Update an application in the following ways: <ul style="list-style-type: none"> • Update the entire application • Add to, replace or delete multiple files in an application • Add a module to an application • Update a module in an application • Delete a module in an application • Add a file to an application • Update a file in an application • Delete a file in an application 	<ul style="list-style-type: none"> • Invoke the AdminApp startApplication command. • Invoke the startApplication method on an ApplicationManager MBean using AdminControl.
Rapid deployment tools See topics under Rapid deployment of J2EE applications .	Briefly, do the following: <ol style="list-style-type: none"> 1. Update your J2EE application files. 2. Set up the rapid deployment environment. 3. Create a free-form project. 4. Launch a rapid deployment session. 5. Drop your updated application files into the free-form project. 	Rapid deployment tools offer the following advantages: <ul style="list-style-type: none"> • You do not need to assemble your J2EE application files prior to deployment. • You do not need to use other installation tools mentioned in this table to deploy the files. 	Use any of the above options to start the application. Clicking Start on the Enterprise applications page is the easiest option.
Hot deployment and dynamic reloading	Briefly, do the following: <ol style="list-style-type: none"> 1. Update your application (.ear), web module (.war), enterprise bean .jar or HTTP plug-in configuration file. 2. Follow instructions in Hot deployment and dynamic reloading to update your file. 	If you are new to WebSphere Application Server, use the administrative console to update applications. That option is easier. Hot deployment and dynamic reloading is more difficult to complete. You must directly manipulate the application or module file on the server where the application is deployed.	Use any of the above options to start the application. Clicking Start on the Enterprise applications page is the easiest option.

You can update .ear, enterprise bean .jar, web module .war, Session Initiation Protocol (SIP) archive (.sar), connector .rar, application client .jar, and any other files used by an installed application.

If the application is updated while it is running, WebSphere Application Server automatically stops the application, updates the application logic and restarts the application. If the application does not start automatically, start it manually using one of the **Starting** options. For more information on the restarting of updated applications, refer to "Fine-grained recycle behavior" in *IBM WebSphere Developer Technical Journal: System management for WebSphere Application Server V6 -- Part 5 Flexible options for updating deployed applications*.

If you update module metadata while an application is running, restarting the application might not be sufficient for the changes to take effect. For example, if you change descriptors in running Java EE 6 applications that use annotations, you must reinstall the application. If you change classes that introduce, remove, or alter class hierarchies within an application, and those changes impact annotated classes, you also must reinstall the application.

Updating enterprise applications with the console

Updating enterprise applications consists of adding a new file or module to an installed Java Platform, Enterprise Edition (Java EE) application, or replacing or removing an installed application, file or module.

Before you begin

Before you update the application files on a server, ensure that the files are assembled in deployable modules.

Next, refer to "Ways to update enterprise application files" on page 322 and decide how to update your application files. You can update enterprise applications or modules using the administrative console, the wsadmin tool, or Java MBean programming. These ways provide similar updating capabilities.

Further, ensure that the updated files can be installed to your deployment targets.

About this task

This topic describes how to update deployed applications or modules using the administrative console.

Procedure

1. Back up the installed application or module.
 - a. Go to the Enterprise applications page of the administrative console.
Click **Applications > Application Types > WebSphere enterprise applications** in the console navigation tree.
 - b. Export the application to an EAR file or export a file in the application.
Select the application you want to export and click **Export** or **Export File**. Exporting preserves the binding information.
2. With the application selected on the Enterprise applications page, click **Update**. The Preparing for application update page is displayed.
3. Under **Specify the EAR, WAR, SAR or JAR module to upload and install**:
 - a. Ensure that **Application to be updated** refers to the application to be updated.
 - b. Under **Application update options**, select the installed application, module, or file that you want to update.
The online help Preparing for application update settings provides detailed information on the options.

Note: You cannot add, remove, or modify a Java Application Programming Interface (API) for XML-Based Web Services (JAX-WS) annotation using the **Replace or add a single file** or **Replace, add, or delete multiple files** update options. These options change a single file or a partial application. If you change a JAX-WS annotation using either of these options, the product does not return an error. However, you might encounter problems deploying annotated web services.

4. If you selected the **Replace the entire application** or **Replace or add a single module** option:
 - a. Click **Next** to display a wizard for updating application files.
 - b. Complete the steps in the update wizard.

This update wizard, which is similar to the installation wizard, provides fields for specifying or editing application binding information. Refer to information on installing applications using the console and on the Preparing for application installation binding settings page 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 Enterprise JavaBeans (EJB) Java Naming and Directory Interface (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.

5. Click **Finish**.
6. If you did not use the Manage modules page of the update wizard, after updating the application, map the installed application or module to servers or clusters.

Use the page accessed from the Enterprise applications page.

- a. Go to the Manage modules page. Click **Applications > Application Types > WebSphere enterprise applications > application_name > Manage modules**.
- b. Specify the application server where you want to install modules contained in your application and click **OK**.

You can deploy Java 2 Platform, Enterprise Edition (J2EE) 1.4 modules to servers on Version 6 or later nodes. You can deploy Java Platform, Enterprise Edition (Java EE) 5 modules to servers on Version 7.x or later nodes. You can deploy Java EE 6 modules to servers on Version 8.x or later nodes.

On multiple-server products, you can also deploy modules to clusters. To deploy J2EE 1.4 modules, the clusters must contain cluster members on Version 6 or later nodes. To deploy Java EE 5 modules, the clusters must contain cluster members on Version 7.x or later nodes. To deploy Java EE 6 modules, the clusters must contain cluster members on Version 8.x or later nodes.

Results

After replacement of a full application, the product uninstalls the old application. After replacement of a module, file or partial application, the product removes the old installed module, file or partial application from the installed application.

What to do next

After the application file or module installs successfully, do the following:

1. If a changed application or module is deployed on a cluster, roll out the changes to all cluster members of the cluster on which the application or module is deployed. Click **Rollout Update** on the Enterprise applications page to propagate the changed configuration on all cluster members of the cluster on which the application or module is deployed. **Rollout Update** sequentially updates the configuration on the nodes that contain cluster members.

best-practices: At the end of the *Installing* messages displayed by the console during application or module installation, click **Manage applications** to go to the Enterprise applications page. Do not save changes to your configuration until after you roll out the changes.

2. Save the changes to your configuration.

On multiple-server products, after you click **Save** the product deletes the old application files and copies new files 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.

3. If needed, restart the application manually so the changes take effect.

If the application is updated while it is running, the product automatically stops the application or only its changed components, updates the application logic, and restarts the stopped application or its components.

If you update module metadata while an application is running, restarting the application might not be sufficient for the changes to take effect. For example, if you change descriptors in running Java EE 6 applications that use annotations, you must reinstall the application. If you change classes that introduce, remove, or alter class hierarchies within an application, and those changes impact annotated classes, you also must reinstall the application.

4. 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, restart the server.

Preparing for application update settings

Use this page to update enterprise applications, modules or files already installed on a server.

To view this administrative console page, do the following:

1. Click **Applications > Application Types > WebSphere enterprise applications**.
2. Select the installed application or module that you want to update.
3. Click **Update**.

Clicking **Update** displays a page that helps you update application files deployed in the cell. You can update the full application, a single module, a single file, or part of the application. If a new file or module has the same relative path as a file or module already existing on the server, the new file or module replaces the existing file or module. If the new file or module does not exist on the server, it is added to the deployed application.

Application to be updated

Specifies the name of the installed (or deployed) application that you selected on the Enterprise applications page.

Replace the entire application

Under **Application update options**, specifies to replace the application already installed on the server with a new (updated) enterprise application `.ear` file.

After selecting this option, do the following:

1. Specify whether the `.ear` file is on a local or remote file system and the full path name of the application. The path provides the location of the updated `.ear` file before installation.

Use **Local file system** if the browser and the updated files or modules are on the same machine, whether or not the server is on that machine too. **Local file system** is available for all update options.

Use **Remote file system** if the application file resides on any node in the current cell context.

On multiple-server installations, using **Remote file system** 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, .sar, or .war files are shown during the browsing.

Also use the **Remote file system** option to specify an application file already residing on the machine running the application server. For example, the field value might be `app_server_install_root/installableApps/test.ear`. If you are installing a stand-alone WAR module, then specify the context root as well.

Tip: During application installation, application files typically are uploaded from a client machine running the browser to the server machine running the administrative console, where they are deployed. In such cases, use the web browser running the administrative console to select 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 **Remote file system** option.

2. If you are installing a stand-alone web application (WAR) or a Session Initiation Protocol (SIP) module (SAR), specify the context root of the WAR or SAR 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`.

3. Click **Next** to display a wizard for updating application files. The update wizard, which is similar to the installation wizard, provides fields for specifying or editing application binding information. Complete the steps in the update wizard as needed.

When the full application is updated, the old application is uninstalled and the new application is installed. When the configuration changes are saved and subsequently synchronized, the application files are expanded on the node where application will run. If the application is running on the node while it is updated, then the application is stopped, application files are updated, and application is started.

Replace or add a single module

Under **Application update options**, specifies to replace a module in or add a module to an installed application.

The module can be a web module (.war file), enterprise bean module (EJB .jar file), SIP module (.sar file), or resource adapter module (connector .rar file).

After selecting this option, specify whether the module is on a local or remote file system and the full path name of the module. The path provides the location of the updated module before installation. For information on **Local file system** and **Remote file system**, refer to the previous description of **Replace the entire application**.

To replace a module, the specified relative path (module URI) must match the path of the module to be updated in the installed application.

To add a new module to the installed application, the specified relative path must *not* match the path of a module in the installed application. The value specifies the desired path for the new module.

If you are installing a stand-alone web or SIP module, specify a value for **Context root**. 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`.

Next, specify whether to show only installation options that require you to supply information or to show all installation options.

After specifying the required information on the module, click **Next** to display a wizard for updating application files. The update wizard, which is similar to the installation wizard, provides fields for specifying or editing module binding information. Complete the steps in the update wizard as needed.

After a single module is added or updated, when configuration changes are saved, the new or updated module is stored in the deployed application in the product configuration repository. When these changes are synchronized with the node, the module is added or updated to the node's file system. If the application is running on the node when the module is added or updated, then one of the following occurs:

- For updates to a web module, the running web module is stopped, web module files are updated, and then the web module is started.
- For module additions, the added module is started on the application servers where the application is running after it is expanded on the node. An application restart is not necessary.
- If the class loader policy for the application is set to `Single` so that all modules share a class loader, then the entire application is stopped and restarted for module level changes.
- If the security provider configured with the product does not support dynamic updates, then the entire application is stopped and restarted for module level changes.
- For all other updates to a module, the entire application is stopped, the module files are updated, then the entire application is started.

Replace or add a single file

Under **Application update options**, specifies to replace a file in or add a file to an installed application.

Use this option to update a file used by the application that is not an `.ear`, `.war`, `.sar`, `.rar` or, in some instances, a `.jar` file. You can use this option to add or update `.jar` files that are not defined as modules in the application. To update an `.ear`, file use the **Replace the entire application** option. To update a `.war` file, `.sar` file, `.rar` file, or `.jar` file that is defined as a module in the application, use the **Replace or add a single module** option.

After selecting this option, specify whether the file is on a local or remote file system and the full path name of the file. The path provides the location of the updated file before installation. For information on **Local file system** and **Remote file system**, refer to the description of **Replace the entire application**.

For the relative path (module URI) , specify a relative path to the file that starts from the root of the `.ear` file. For example, if the file is located at `com/company/greeting.class` in module `hello.jar`, specify a relative path of `hello.jar`.

To replace a file, the relative path must match the relative path of the file to be updated in the installed application.

To add a new file to the installed application, the specified relative path must *not* match the relative path of an already existing file in the installed application. The value specifies the desired path for the new file.

After you specify the file system and relative paths, click **Next**.

After a single file is added or updated, when configuration changes are saved, the new or updated file is stored in the deployed application in the product configuration repository. When these changes are synchronized with the node, the file is added or updated to the node's file system. If the application is running on the node when the file is added or updated, then one of the following occurs:

- When files are added at application metadata scope (META-INF directory) or updated at any application scope or in non-web modules, the entire application is stopped, the file is added or updated, and then the entire application is restarted.
- When files are added at application non-metadata scope (outside of META-INF directory but not in any module), the changes are saved in the file system without restarting the running application.
- When files are added or updated to web module metadata (META-INF or WEB-INF directory), the running web module is stopped, the web module file is added or updated, and then the web module is started.

- For all other files in web modules, the file is added or updated on the node's file system without stopping the application or any of its components.

Replace, add, or delete multiple files

Under **Application update options**, specifies to update multiple files of an installed application by uploading a compressed file. Depending on the contents of the compressed file, a single use of this option can replace files in, add new files to, and delete files from the installed application. Each entry in the compressed file is treated as a single file and the path of the file from the root of the compressed file is treated as the relative path of the file in the installed application.

After selecting this option, specify whether the compressed file is on a local or remote file system and the full path name of the compressed file. You will likely use **Local file system** because you are uploading a compressed file and remote browsing only works for .ear, .sar, .war or .jar files. Specify a valid compressed file format such as .zip or .gzip. The path provides the location of the compressed file before installation. This option unzips the compressed file into the installed application directory.

Use **Local file system** if the browser and the updated files or modules are on the same machine, whether or not the server is on that machine too. **Local file system** is available for all update options.

To replace a file, a file in the compressed file must have the same relative path as the file to be updated in the installed application.

To add a new file to the installed application, a file in the compressed file must have a different relative path than the files in the installed application.

The relative path of a file in the installed application is formed by concatenation of the relative path of the module (if the file is inside a module) and the relative path of the file from the root of the module separated by /.

To remove a file from the installed application, specify metadata in the compressed file using a file named META-INF/ibm-partialapp-delete.props at any archive scope. The ibm-partialapp-delete.props file must be an ASCII file that lists files to be deleted in that archive with one entry for each line. The entry can contain a string pattern such as a regular expression that identifies multiple files. The file paths for the files to be deleted must be relative to the archive path that has the META-INF/ibm-partialapp-delete.props file.

Level of files to delete	Metadata .props file to include in compressed file
Application	<p>Include META-INF/ibm-partialapp-delete.props in the compressed file. In the metadata .props file, list files to be deleted. File paths are relative to the location of the META-INF/ibm-partialapp-delete.props file.</p> <p>For example, to delete a file named utils/config.xml from the root of the my.ear file, include the line utils/config.xml in the META-INF/ibm-partialapp-delete.props file.</p>
Module	<p>Include <i>module_uri</i>/META-INF/ibm-partialapp-delete.props in the compressed file.</p> <p>To delete one file from a module, include the file path relative to the module in the metadata .props file. For example, to delete a/b/c.jsp from the my.jar module, include a/b/c.jsp in my.jar/META-INF/ibm-partialapp-delete.props file in the compressed file.</p> <p>To delete multiple files within a module, list the files to be deleted in the metadata .props file with one entry on each line. For example, to delete all JavaServer Pages (.jsp files) from the my.war file, include the line *.jsp in the my.war/META-INF/ibm-partialapp-delete.props file. The line uses a regular expression, *.jsp, to identify all .jsp files in my.war.</p>

You can use a single partial application file to add, delete and update multiple files.

After you specify a file system path, click **Next**.

After a partial application update, when configuration changes are saved, the new or updated application file is stored in the deployed application in the WebSphere Application Server configuration repository. When these changes are synchronized with the node, the files are added or updated to the node's file system. Because the partial application option updates multiple files, the application components that are restarted are determined using individual files in the partial application.

An example of entries in a partial application compressed file follows:

```
util.jar
META-INF/ibm-partialapp-delete.props
foo.jar/com/mycomp/xyz.class
xyz.war/welcome.jsp
xyz.war/WEB-INF/web.xml
webmod.war/META-INF/ibm-partialapp-delete.props
```

For this example, the META-INF/ibm-partialapp-delete.props file contains the *.dat and tools/test.jar files. The webmod.war/META-INF/ibm-partialapp-delete.props file contains the com/test/*.jsp and WEB-INF/test.xmi files.

The partial application update option does the following:

- Adds or replaces util.jar in the deployed application.
- Adds or replaces com/mycomp/xyz.class inside the foo.jar file of the deployed application.
- Deletes *.dat files from the application, but not from any modules.
- Deletes tools/test.jar from the application.
- Adds or replaces welcome.jsp inside the xyz.war module of the deployed application.
- Replaces WEB-INF/web.xml inside the xyz.war module of the deployed application.
- Deletes com/test/*.jsp from the webmod.war module.
- Deletes WEB-INF/test.xmi from the webmod.war module.

Escape regular expression metacharacters in the META-INF/ibm-partialapp-delete.props file. For example, to delete inner classes for a class named Abc, use the regular expression Abc\<\$.* where \$ is a regular expression metacharacter that is escaped with a backslash (\).

A META-INF/ibm-partialapp-delete.props file might contain the following text:

```
*.dat

webmod.war/META-INF/ibm-partialapp-delete.props:
com/test/*.jsp
WEB-INF/test.xmi
```

Hot deployment and dynamic reloading

You can make various changes to applications and their modules without having to stop the server and start it again. Making these types of changes is known as *hot deployment and dynamic reloading*.

Before you begin

The following note applies to the xmi file references in this topic:

Note: For IBM extension and binding files, the .xmi or .xml file name extension is different depending on whether you are using a pre-Java EE 5 application or module or a Java EE 5 or later application or module. An IBM extension or binding file is named ibm-*-ext.xmi or ibm-*-bnd.xmi where * is the type of extension or binding file such as app, application, ejb-jar, or web. The following conditions apply:

- For an application or module that uses a Java EE version prior to version 5, the file extension must be .xmi.

- For an application or module that uses Java EE 5 or later, the file extension must be .xml. If .xmi files are included with the application or module, the product ignores the .xmi files.

However, a Java EE 5 or later module can exist within an application that includes pre-Java EE 5 files and uses the .xmi file name extension.

The `ibm-webservices-ext.xmi`, `ibm-webservices-bnd.xmi`, `ibm-webservicesclient-bnd.xmi`, `ibm-webservicesclient-ext.xmi`, and `ibm-portlet-ext.xmi` files continue to use the .xmi file extensions.

This topic assumes that your application files are deployed on a server and you want to upgrade the files.

See “Ways to update enterprise application files” on page 322 and determine whether hot deployment is the appropriate way for you to update your application files. Other ways are easier and hot deployment is appropriate only for experienced users.

Do not use hot deployment if you intend to export your application, generate a plug-in based on the application configuration, or perform other application management in the future. Changes that you make to your application files using hot deployment are not recognized by administrative console or wsadmin application management functions. Those functions recognize only the application files that administrative programs such as the console or wsadmin present during application installation, update or other management functions. The application management functions do not recognize files changed by hot deployment.

Important: Do not use hot deployment to update components in a production deployment manager managed cell. Hot deployment is well-suited for development and testing, but poses unacceptable risks to production environments. Full or partial resynchronization might erase hot deployed components. Also, running the `restoreconfig` command might overwrite changes made to expanded application files. Further, hot deployed components are not migrated between versions of WebSphere Application Server. To add new components or modules to an enterprise application, reassemble the application EAR file so it has the new components or modules and then redeploy the EAR file.

About this task

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

As opposed to the changes made to a deployed application described in “Updating enterprise application files” on page 320, changes made using hot deployment or dynamic reloading do not use the administrative console or a wsadmin scripting command. You must directly manipulate the application files on the server where the application is deployed.

If the application you are updating is deployed on a server that has its application class loader policy set to `Single`, you might not be able to dynamically reload your application. At minimum, you must restart the server after updating your application.

Procedure

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, *app_server_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.

Important: You can have *useMetadataFromBinaries=true*, change an extracted copy of your application using hot deployment, and have the changes take effect at run time by following the procedure in this topic. However, changes that you make to your application files using hot deployment are not recognized by console or wsadmin application management functions. Those functions recognize only the original application files and not the files changed by hot deployment. Do not use hot deployment if you intend to export your application, generate a plug-in based on the application configuration, or perform other application management in the future. Hot deployment enables you to quickly change application files; it does not support the full management lifecycle of an application.

For the remainder of this information, *metadata_root* represents the location of the metadata files for the specified application or module.

3. Required: If you are running WebSphere Application Server on a group of machines using WebSphere Application Server, Network Deployment and you are changing an application on a particular node, disable automatic synchronization.
 - a. Click **System Administration > Node agents > node_agent_name > File synchronization service** in the console navigation tree.
 - b. On the File synchronization service page, clear the check box for **Automatic synchronization** and click **OK**.

When you run WebSphere Application Server on a group of machines using WebSphere Application Server, 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 WebSphere Application Server, 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 classes when application files are updated** and **Polling interval for updated files** on the settings page for your application's class loader.

If reloading of classes is enabled and the polling interval is greater than zero (0), the application files are reloaded after the application is updated. For JavaServer Pages (JSP) files in a web module, a web container reloads JSP files only when the IBM extension `jspReloadingEnabled` in the `jspAttributes` of the `ibm-web-ext.xmi` file is set to `true`. You can set `jspReloadingEnabled` 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 or stopping enterprise applications” on page 316 provides information on using the administrative console to start, stop, or restart an application.
7. If you disabled automatic synchronization in step 3, enable automatic synchronization again:
 - a. Return to the File synchronization service page.
 - b. Select **Automatic synchronization**.
 - c. Click **OK**.

Results

The application files are updated on the server.

Because you directly manipulated the application files on the server, you might not be able to later use the administrative console or a `wsadmin` scripting command to work with the files. For example, if you try exporting a manually changed application using **Export** on an Enterprise applications console page, your manual changes to an application in the `installedApps` directory are not exported. To export those changes, you must copy and move the application files manually.

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.

About this task

The following note applies to the `xmi` file references in this topic:

Note: For IBM extension and binding files, the `.xmi` or `.xml` file name extension is different depending on whether you are using a pre-Java EE 5 application or module or a Java EE 5 or later application or module. An IBM extension or binding file is named `ibm-*-ext.xmi` or `ibm-*-bnd.xmi` where `*` is the type of extension or binding file such as `app`, `application`, `ejb-jar`, or `web`. The following conditions apply:

- For an application or module that uses a Java EE version prior to version 5, the file extension must be `.xmi`.
- For an application or module that uses Java EE 5 or later, the file extension must be `.xml`. If `.xmi` files are included with the application or module, the product ignores the `.xmi` files.

However, a Java EE 5 or later module can exist within an application that includes pre-Java EE 5 files and uses the `.xmi` file name extension.

The `ibm-webservices-ext.xmi`, `ibm-webservices-bnd.xmi`, `ibm-webservicesclient-bnd.xmi`, `ibm-webservicesclient-ext.xmi`, and `ibm-portlet-ext.xmi` files continue to use the `.xmi` file extensions.

There are several changes that you can make to deployed application files without stopping the server and starting it again.

Important: See “Ways to update enterprise application files” on page 322 and determine whether hot deployment is the appropriate way for you to update your application files. Other ways are easier and hot deployment is appropriate only for experienced users. You can use the update wizard of the administrative console to make the changes without having to stop and restart the server.

The following table lists the changes that you can make by manipulating an application file on the server where the application is deployed. The table also states whether you use hot deployment or dynamic reloading to make the changes.

Table 37. Available changes to deployed application files. Available changes using hot deployment or dynamic reloading.

Change	Hot deployment	Dynamic reloading
Update an existing application on a running server by providing a new EAR file.	Yes	Yes
Add a new application to a running server.	Yes	No
Remove an existing application from a running server.	Yes	No
Change or add files to existing EJB or web modules.	Yes	No
Change the <code>application.xml</code> file for an application.	Not applicable	Yes
Change the <code>ibm-app-ext.xmi</code> file for an application.	Not applicable	Yes
Change the <code>ibm-app-bnd.xmi</code> file for an application.	Not applicable	Yes
Change a non-module Jar file contained in the EAR file.	Yes	Yes

Procedure

- Update an existing application on a running server by 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.
- Add a new application to a running server.
Install an application using the administrative console or the `wsadmin install` command.
- Remove 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 use the `stopApplication` attribute of the `AdminControl` object with the `wsadmin` tool and then run the `uninstall` command.
- Change or add 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 use the `startApplication` and `stopApplication` attributes of the `AdminControl` object with the `wsadmin` tool.
- Change 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 use the `startApplication` and `stopApplication` attributes of the `AdminControl` object with the `wsadmin` tool.

- Change 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 use the `startApplication` and `stopApplication` attributes of the `AdminControl` object with the `wsadmin` tool.

- Change 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 use the `startApplication` and `stopApplication` attributes of the `AdminControl` object with the `wsadmin` tool.

- Change 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 use the `startApplication` and `stopApplication` attributes of the `AdminControl` object with the `wsadmin` tool.

If automatic reloading is enabled, you do not need to take further action. Automatic reloading will detect the change.

Changing or adding WAR files

You can change web application archives (WAR files) on application servers without having to stop the server and start it again.

About this task

The following note applies to the file references with a `.xmi` extension in this topic:

Note: For IBM extension and binding files, the `.xmi` or `.xml` file name extension is different depending on whether you are using a pre-Java EE 5 application or module or a Java EE 5 or later application or module. An IBM extension or binding file is named `ibm-*-ext.xmi` or `ibm-*-bnd.xmi` where `*` is the type of extension or binding file such as `app`, `application`, `ejb-jar`, or `web`. The following conditions apply:

- For an application or module that uses a Java EE version prior to version 5, the file extension must be `.xmi`.
- For an application or module that uses Java EE 5 or later, the file extension must be `.xml`. If `.xmi` files are included with the application or module, the product ignores the `.xmi` files.

However, a Java EE 5 or later module can exist within an application that includes pre-Java EE 5 files and uses the `.xmi` file name extension.

The `ibm-webservices-ext.xmi`, `ibm-webservices-bnd.xmi`, `ibm-webservicesclient-bnd.xmi`, `ibm-webservicesclient-ext.xmi`, and `ibm-portlet-ext.xmi` files continue to use the `.xmi` file extensions.

There are several changes that you can make to WAR files without stopping the server and starting it again.

Important: See “Ways to update enterprise application files” on page 322 and determine whether hot deployment is the appropriate way for you to update your WAR files. Other ways are easier and hot deployment is appropriate only for experienced users. You can use the update wizard of the administrative console to make the changes without having to stop and restart the server.

The following table lists the changes that you can make by manipulating a WAR file on the server where the application is deployed. The table also states whether you use hot deployment or dynamic reloading to make the changes.

Table 38. Available changes to deployed WAR files. Available changes using hot deployment or dynamic reloading.

Change	Hot deployment	Dynamic reloading
Change an existing JavaServer Pages (JSP) file.	Not applicable	Yes
Add a new JSP file to an existing application.	Yes	Yes
Change an existing servlet class by editing and recompiling.	Not applicable	Yes
Change a dependent class of an existing servlet class.	Not applicable	Yes
Add a new servlet using the Invoker (Serve Servlets by class name) facility or add a dependent class to an existing application.	Yes	Not applicable
Add a new servlet, including a new definition of the servlet in the web.xml deployment descriptor for the application.	Yes	Not applicable
Change the web.xml file of a WAR file.	Yes	Yes
Change the ibm-web-ext.xml file of a WAR file.	Not applicable	Yes
Change the ibm-web-bnd.xml file of a WAR file.	Not applicable	Yes

Procedure

- Change an existing JavaServer Pages (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.
- Add 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.
- Change an existing servlet class by 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 use the `startApplication` and `stopApplication` attributes of the `AdminControl` object with the `wsadmin` tool.
If automatic reloading is enabled, you do not need to take further action. Automatic reloading will detect the change.
- Change 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 use the `startApplication` and `stopApplication` attributes of the `AdminControl` object with the `wsadmin` tool.
If automatic reloading is enabled, you do not need to take further action. Automatic reloading will detect the change.
- Add a new servlet using the Invoker (Serve Servlets by class name) facility or add 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 use the `startApplication` and `stopApplication` attributes of the `AdminControl` object with the `wsadmin` tool.

If automatic reloading is enabled, you do not need to take further action. Automatic reloading will detect the change.

- Add 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 use the `startApplication` and `stopApplication` attributes of the `AdminControl` object with the `wsadmin` tool. After the application restarts, the new servlet is available for service.
- Change 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 use the `startApplication` and `stopApplication` attributes of the `AdminControl` object with the `wsadmin` tool.
- Change 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 no longer automatically detects 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.

Note: The `reloadInterval` and `reloadEnabled` attributes of the IBM deployment descriptor extensions, including both the WAR file extension `WEB-INF/ibm-web-ext.xmi` and the application extension `META-INF/ibm-application-ext.xmi` are deprecated.

- Change 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 use the `startApplication` and `stopApplication` attributes of the `AdminControl` object with the `wsadmin` tool.

Changing or adding EJB JAR files

You can change enterprise bean (EJB) Java archive (JAR) files on application servers without having to stop the server and start it again.

About this task

The following note applies to the file references with `.xmi` extensions in this topic:

Note: For IBM extension and binding files, the .xmi or .xml file name extension is different depending on whether you are using a pre-Java EE 5 application or module or a Java EE 5 or later application or module. An IBM extension or binding file is named `ibm-*-ext.xmi` or `ibm-*-bnd.xmi` where * is the type of extension or binding file such as `app`, `application`, `ejb-jar`, or `web`. The following conditions apply:

- For an application or module that uses a Java EE version prior to version 5, the file extension must be `.xmi`.
- For an application or module that uses Java EE 5 or later, the file extension must be `.xml`. If `.xmi` files are included with the application or module, the product ignores the `.xmi` files.

However, a Java EE 5 or later module can exist within an application that includes pre-Java EE 5 files and uses the `.xmi` file name extension.

The `ibm-webservices-ext.xmi`, `ibm-webservices-bnd.xmi`, `ibm-webservicesclient-bnd.xmi`, `ibm-webservicesclient-ext.xmi`, and `ibm-portlet-ext.xmi` files continue to use the `.xmi` file extensions.

There are several changes that you can make to EJB JAR files without stopping the server and starting it again.

Important: See “Ways to update enterprise application files” on page 322 and determine whether hot deployment is the appropriate way for you to update your EJB JAR files. Other ways are easier and hot deployment is appropriate only for experienced users. You can use the update wizard of the administrative console to make the changes without having to stop and restart the server.

The following table lists the changes that you can make to EJB JAR files by manipulating an EJB file on the server where the application is deployed. The table also states whether you use hot deployment or dynamic reloading to make the changes.

Table 39. Available changes to EJB JAR files. Available changes using hot deployment or dynamic reloading.

Change	Hot deployment	Dynamic reloading
Change the <code>ejb-jar.xml</code> file of an EJB JAR file.	Not applicable	Yes
Change the <code>ibm-ejb-jar-ext.xmi</code> or <code>ibm-ejb-jar-bnd.xmi</code> file of an EJB JAR file.	Not applicable	Yes
Change the <code>Table.ddl</code> file for an EJB JAR file.	Not applicable	Not applicable
Change the <code>Map.mapxmi</code> or <code>Schema.dbxmi</code> file for an EJB JAR file.	Not applicable	Yes
Update the implementation class for an EJB file or a dependent class of the implementation class for an EJB file.	Not applicable	Yes
Update the Home/Remote interface class for an EJB file.	Not applicable	Yes
Add a new EJB file to an existing EJB JAR file.	Yes	Yes

Procedure

- Change 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 use the `startApplication` and `stopApplication` attributes of the `AdminControl` object with the `wsadmin` tool.
- Change the `ibm-ejb-jar-ext.xmi` or `ibm-ejb-jar-bnd.xmi` 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 use the `startApplication` and `stopApplication` attributes of the `AdminControl` object with the `wsadmin` tool.

- Change 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.
- Change the Map.mapxmi or Schema.dbxmi file for an EJB JAR file.
 1. Change the Map.mapxmi or Schema.dbxmi 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 use the startApplication and stopApplication attributes of the AdminControl object with the wsadmin tool.
- Update 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 use the startApplication and stopApplication attributes of the AdminControl object with the wsadmin tool.
- Update 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 use the startApplication and stopApplication attributes of the AdminControl object with the wsadmin tool.
- Add 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 use the startApplication and stopApplication attributes of the AdminControl object with the wsadmin tool.

Changing the HTTP plug-in configuration

You can change the HTTP plug-in configuration without having to stop the server and start it again.

About this task

There are several change that you can make to the HTTP plug-in configuration without stopping the server and starting it again.

Important: See “Ways to update enterprise application files” on page 322 and determine whether hot deployment is the appropriate way for you to update your HTTP plug-in configuration. Other ways are easier and hot deployment is appropriate only for experienced users.

The following table lists the changes that you can make to the HTTP plug-in configuration. The table also states whether you use hot deployment or dynamic reloading to make the changes.

Table 40. Available changes to HTTP plug-in configuration files. Available changes using hot deployment or dynamic reloading.

Change	Hot deployment	Dynamic reloading
Change the <code>application.xml</code> file to change the context root of a web application archive (WAR file).	Yes	No
Change the <code>web.xml</code> file to add, remove, or modify a servlet mapping.	Yes	Yes
Change the <code>server.xml</code> file to add, remove, or modify an HTTP transport or change the <code>virtualhost.xml</code> file to add or remove a virtual host or to add, remove, or modify a virtual host alias.	Yes	Yes

Note: The following steps reference the `GenPluginCfg.bat/sh` script. You must delete the `plugin-cfg.xml` file in the `profile_root/config/cells` directory before you use the `GenPluginCfg.bat/sh` script. Otherwise, configuration changes do not persist to the `plugin-cfg.xml` file.

Procedure

- Change the `application.xml` file to change the context root of a WAR file.
 1. Change the `application.xml` file.
 2. If the plug-in configuration property **Automatically propagate plug-in configuration file** is selected for this plug-in, it is automatically regenerated whenever the `application.xml` file changes.
See documentation on the web server plug-in properties for information on how to set this property. You can also run the `GenPluginCfg.bat/sh` script, or issue a `wsadmin` command to regenerate the plug-in configuration file.
- Change the `web.xml` file to add, remove, or modify a servlet mapping.
 1. Change the `web.xml` file.
 2. If the plug-in configuration property **Automatically propagate plug-in configuration file** is selected for this plug-in, it is automatically regenerated whenever the `web.xml` file changes.
See documentation on the web server plug-in properties for information on how to set this property. You can also run the `GenPluginCfg.bat/sh` script, or issue a `wsadmin` command to regenerate the plug-in configuration file.
If the web application has file serving enabled or has a servlet mapping of `/`, the plug-in configuration does not have to be regenerated. In all other cases a regeneration is required.
- 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.
 1. Change the `server.xml` file or the `virtualhost.xml` file.
 2. If the plug-in configuration property **Automatically propagate plug-in configuration file** is selected for this plug-in, it is automatically regenerated whenever the `server.xml` file changes.
See documentation on the web server plug-in properties for information on how to set this property. You can also run the `GenPluginCfg.bat/sh` script, or issue a `wsadmin` command to regenerate the plug-in configuration file.

Resolving application configuration conflicts

In a shared environment with multiple administrative users, it is possible that different administrative users might attempt concurrent updates of the same WebSphere Application Server configuration documents. The following information should help you detect and deal with any exceptions that might occur if multiple administrative users attempt to concurrently update the same configuration documents.

Whenever you log into an administrative client using either the administrative console or `wsadmin` tool, a unique workspace session is created to track any configuration changes that are made. For each workspace session, a temporary workspace directory is associated with each workspace session. This directory is used to store all of the configuration files that you change during a login session. The files in

this directory are initially extracted from the cell configuration repository, and your changes only exist in the workspace copies of these files until a save occurs. When a save occurs, the configuration management runtime makes sure that the configuration files you changed have not already been modified and saved by another user, and then copies the changed files from your workspace directory back to the master repository,

As long as different workspace sessions modify different configuration files, there is no save conflict. However if multiple workspace sessions modify one or more of the same configuration files, a save conflict occurs, and only the first workspace session changes are reflected in the configuration repository for the cell. When subsequent users attempt to save the changes to the same configuration files, they receive a save conflict exception

With regard to application deployment, a special provision is made for supporting the concurrent deployment of applications. Application deployment initiated by different workspace sessions can modify the same `serverindex.xml` file. In this situation, the configuration management runtime employs a merge algorithm for the `serverindex.xml` file which supports:

- Concurrent deployment of different applications to different application servers or to different application server clusters.
- Concurrent deployment of different applications to the same application server or to the same application server cluster.

Typically, concurrent or parallel application deployment scenarios are safely executed without any additional effort on the part of the system administrator. However, there are application deployment scenarios that the `serverindex.xml` merge algorithm does not handle. For example, the `serverindex.xml` merge algorithm does not handle situations where file changes are concurrently deployed to the same application, the same cluster or the same server. The merge algorithm also does not handle configuration conflicts that arise during other concurrent administration activities that involve more than application deployment.

Obtaining the requisite object references and constructing the parameter list

There are some simple measures you can take to ensure against configuration conflicts, and to resolve any conflicts when using the `wsadmin` tool. The `wsadmin` commands that are employed for configuration conflict detection and resolution rely on obtaining a reference to the `ConfigService` MBean, and then invoking the `getConflictDocuments` method provided by that MBean to determine if users have made conflicting changes to a file during their workspace session. See the Javadoc for the `ConfigService` MBean for more information about this Mbean.

The following code example illustrates how to obtain the requisite object references and construct the parameter list that is required to invoke the `getConflictDocuments` method that the `ConfigService` MBean provides:

```
// get ConfigService MBean reference
wsadmin>cs = AdminControl.queryNames('WebSphere:*,type=ConfigService')

// obtain ObjectName for ConfigService MBean
wsadmin>import javax.management as mgmt
wsadmin>csName=mgmt.ObjectName(cs)

// get session object for the current administrative user session
wsadmin>session=AdminConfig.getCurrentSession()

// manipulate and prepare the administrative session object and
// MBean operation arguments for use
wsadmin>from com.ibm.websphere.management import Session
wsadmin>from jarray import array
```

```
wsadmin>parms=array([session], java.lang.Object)
wsadmin>ptype=array(['com.ibm.websphere.management.Session'], java.lang.String)
```

After the variables and parameter list is initialized, the `getConflictDocuments` method is invoked. If there are no conflicts, the method returns the following message:

```
// invoke MBean getConflictDocuments method to obtain a list of any document conflicts
wsadmin>AdminControl.invoke_jmx(csName,'getConflictDocuments', parms, ptype)
{}
wsadmin>
```

If configuration conflicts exist because of changes another user made during a work session, the method returns a message, similar to the following message, that lists the XML files that have changed:

```
Listing 3
wsadmin>AdminControl.invoke_jmx(csName,'getConflictDocuments', parms, ptype)
{['cells/cell_name/nodes/node_name/serverindex.xml','cells/cell_name/applications/
DefaultApplication.ear.ear/deltas/DefaultApplication.ear/delta-1278791909117',
... <list abbreviated> ...}
wsadmin>
```

In this situation, you can issue the `AdminConfig.reset()` command to discard the changes that you made since the last `AdminConfig.save()` command was issued:

```
wsadmin>AdminConfig.reset()
```

Even if you call the `getConflictDocuments` method prior to saving your changes, and see that there are no conflict documents, there is no guarantee that a save will succeed, even if you immediately issue the `AdminConfig.save()` command because some other session might have modified the same configuration files between when you call the `getConflictDocuments` method, and when you issue the `AdminConfig.save()` command.

When a save to the master repository is unsuccessful, you get a `ConfigServiceException` exception that is similar to the following exception:

```
WASX7015E: Exception running command: "AdminConfig.save()"; exception information:
com.ibm.websphere.management.exception.ConfigServiceException
java.security.PrivilegedActionException:
java.security.PrivilegedActionException:
com.ibm.ws.sm.workspace.WorkSpaceException: RepositoryException
```

If you receive a save conflict exception:

1. Use the `getConflictDocuments` method to determine which configuration files were already saved by another user.
2. Issue the `AdminConfig.reset()` command to discard your changes.
3. After discarding these changes, you can reapply your changes to the appropriate configuration files, and issue the `AdminConfig.save()` command to save these changes.

The subsequent attempt to save your changes is likely to succeed because it is not typical to experience multiple save conflicts during the same session. However, if this subsequent save attempt is not successful, repeat the preceding actions, and save your changes again.

Exporting enterprise applications

You can export an enterprise application to a location of your choice.

Before you begin

This topic assumes that you have installed an enterprise application on a server and that you want to export the application.

About this task

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 product.

To export applications, use the **Export** button on the Enterprise applications page. Using **Export** produces an enhanced enterprise archive (EAR) file that contains the application as well as the deployment configuration. The *deployment configuration* consists of the `deployment.xml` and other configuration files that control the application behavior on a deployment target.

Procedure

1. Click **Applications > Application Types > WebSphere enterprise applications** in the console navigation tree to access the Enterprise applications page.
2. Select 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.
5. Click **Back** to return to the Enterprise applications page.

Results

The file containing binding information is exported to the specified node and directory, and has the name *enterprise_application_name.ear*.

Using the **Export** button to export applications does not export any manual changes that were made to applications in the `installedApps` directory. To export those changes, you must copy and move the application files manually.

What to do next

You can edit your exported enhanced EAR file and then reinstall it. By default, installation expands an EAR file in the *profile_root/installedApps/cell_name* directory. If you specified the `$(CELL)` variable for **Directory to install application** on the Select installation options panel of the application installation wizard when you first installed the application, the *cell_name* directory is the current cell name.

To reinstall the enhanced EAR file, do either of the following:

- Use the **Update** operation available from the Enterprise applications page to upgrade the existing application installation.

The **Update** operation adds the application files to the *profile_root/installedApps/cell_name* directory, where *cell_name* is the current cell name or the name of the cell that you specified for **Directory to install application** when you first installed the application on a deployment target. The **Directory to install application** setting is on the Select installation options panel of the application installation wizard. If you specified the `$(CELL)` variable for **Directory to install application** when you first installed the application, the *cell_name* directory is the current cell name.

- Use the **Applications > New application > New Enterprise Application** operation to install the exported EAR file.

If you specified the `$(CELL)` variable for **Directory to install application** when you first installed the application, the *cell_name* directory is the current cell name. That is, if the file is originally installed on Cell1 with `$(CELL)` variable in the destination directory and you reinstall the enhanced EAR file on Cell2, the *cell_name* directory is Cell2, the current cell name.

If the `$(CELL)` variable was not specified for the first installation, using **New Enterprise Application** to reinstall an enhanced EAR file installs the application in the *cell_name* directory of the exported application. That is, if the application is originally installed on and exported from Cell1 and you reinstall

the enhanced EAR file on Cell2, the *cell_name* directory is Ce111. The enhanced EAR file expands in the Ce111 directory even though the current cell name is Cell2. By default, the application destination directory contains Ce111 in its path because the `deployment.xml` file in the exported application has Cell1 in it.

If you exported the application from Cell1 and did not specify the `$(CELL)` variable when first installing the application, and you want to install the enhanced EAR file on a different cell, deselect **Process embedded configuration** on the Select installation options panel of the application installation wizard to expand the enhanced EAR file in the current cell name directory, which is not Ce111.

Exporting enterprise application files

You can export individual files of a Java Platform, Enterprise Edition (Java EE) application or module.

Before you begin

This topic assumes that you have installed an application or module on a server and that you want to export a file in the application or module.

About this task

Exporting a file in a deployed application or module downloads the file to a location of your choice.

To export a file using the administrative console, use **Export File**.

To export an entire application, use **Export**. For information on **Export**, see “Exporting enterprise applications” on page 342. The exported enterprise archive (EAR) file contains application configuration data as well as the application.

Procedure

1. Click **Applications > Application Types > WebSphere 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 File**. A drop-down list of exportable files is displayed.
3. Select a file from the list and click **Export**. A dialog in which you select a target location is displayed. If the browser does not prompt for a location to store the file, click **File > Save as** and specify a location to save the file that is shown in the browser.
4. Specify the location to which to download the file.

Results

The file is downloaded to the specified location.

What to do next

Click **Back** to return to the Enterprise applications page.

Exporting DDL files

You can export data definition language (DDL) files in the enterprise bean (EJB) modules of an application.

About this task

Exporting DDL (Table.ddl) files in the EJB modules of an application downloads the DDL files to a location of your choice.

Procedure

1. Click **Applications > Application Types > WebSphere 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 DDL**. If the application has no DDL files in any of its EJB modules, then the message *No DDL files were found* is displayed at the top of the page. If the application has DDL files in its EJB modules, then a page listing DDL files in the format *application_name.ear/_module.jar_Table.ddl* is displayed.
3. Click on a file in the list and specify the location to which to download the file.

Tip: For Firefox browsers, right-click the file name, select **Save Link As**, and specify the location to which to download the file.

Mozilla browsers might display the contents of the Table.ddl file instead of saving the file to disk. To save the file, edit the **Helper Application** preference settings of the Mozilla browser by adding a new type for DDL and specifying that you want to save DDL files to disk. That is, set MIME type = `ddl` and Extension = `ddl`.

Results

The product downloads the DDL file to the specified location.

Uninstalling enterprise applications using the console

After an application no longer is needed, you can uninstall it.

Before you begin

This topic assumes that you have installed an enterprise application on a server and that you want to delete the application from the server.

About this task

Uninstalling an application deletes the application from the product configuration repository and deletes the application binaries from the file system of all nodes where the application modules are installed.

Procedure

1. Click **Applications > Application Types > WebSphere enterprise applications** in the administrative console navigation tree to access the Enterprise applications page.
2. If you need to retain a copy of the application, back up the application.

- a. Select the application to uninstall.
- b. Click **Export**.

The product exports the application to an enterprise application (`.ear`) file, preserving the binding information.

3. Uninstall the application.
 - a. Select the application to uninstall.
 - b. Click **Uninstall**.
 - c. On the Uninstall application page, click **OK**.

4. Save changes made to the administrative configuration.

Results

On single-server products, application binaries are deleted after you save the changes.

On multiple-server products, application binaries are deleted when configuration changes on the deployment manager synchronize with configurations for individual nodes.

Removing enterprise files

After a file is no longer needed, you can remove the file from a Java Platform, Enterprise Edition (Java EE) application or module deployed on a server.

Before you begin

This topic assumes that you have installed an application or module file on a server and that you want to delete the file.

About this task

Removing a file deletes the file from the product configuration repository and deletes the file from the file system of all nodes where the file is installed.

You can use the administrative console to remove a file from an application or module.

Procedure

- Remove a file from an application.
 1. Go to the Enterprise applications page.
Click **Applications > Application Types > WebSphere enterprise applications** in the console navigation tree.
 2. Select the application that contains a file you want removed.
 3. Click **Remove File**. The Remove a file page is displayed
 4. Select the URI of the file that you want removed from the application.
 5. Back up the application.
Under **Export before removing file**, select the application name and then specify the location to which you want the file exported.
 6. Click **OK** to remove the file.
- Remove a file from a module.
 1. Go to the Manage modules page.
Click **Applications > Application Types > WebSphere enterprise applications > *application_name* > Manage modules** in the console navigation tree.
 2. Select the module from which you want to delete a file.
 3. Click **Remove File**. The Remove a file from a module page is displayed.
 4. Select the URI of the file that you want removed from the module.
 5. Back up the application.
Under **Export before removing file**, select the application name and then specify the location to which you want the file exported.
 6. Click **OK** to remove the file.

Results

The file is exported to the designated location and removed from the application or module. The application or stand-alone Web module that had a file removed is restarted so the changes take effect.

What to do next

If the application or module is deployed on a cluster and you have no more configuration changes to make, click **Rollout Update** on the Enterprise applications page to propagate the changed configuration on all cluster members of the cluster on which the application or module is deployed. **Rollout Update** sequentially updates the configuration on the nodes that contain cluster members.

Save the changes to your administrative configuration.

On single-server products, application binaries are deleted after you save the changes.

On multiple-server products, application binaries are deleted when configuration changes on the deployment manager synchronize with configurations for individual nodes.

Uninstalling enterprise application files by dragging them from a monitored directory

After an application file no longer is needed, you can uninstall it by removing its enterprise archive (EAR), web application archive (WAR), Java archive (JAR), or Session Initiation Protocol (SIP) archive (SAR) from a monitored directory. An enterprise application file must conform to the Java Platform, Enterprise Edition (Java EE) specification.

Before you begin

This topic assumes that you previously installed an enterprise application or module on a server or cluster by dragging or copying it to a monitored directory and that you want to delete it from the server or cluster. The EAR, JAR, WAR, or SAR file exists in a monitored subdirectory:

- *app_server_root/profiles/default/monitoredDeployableApps*

For deployment managers on z/OS operating systems, the monitored directory path is as follows:

- For application servers, the directory path is *app_server_root/AppServer/profiles/default/monitoredDeployableApps/servers/server_name*.
- For managed (federated) nodes, the directory path is *app_server_root/DeploymentManager/profiles/default/monitoredDeployableApps/nodes/node_name/servers/server_name*.
- For clusters, the directory path is *app_server_root/DeploymentManager/profiles/default/monitoredDeployableApps/clusters/cluster_name*.

Restriction: Installing an EAR, JAR, WAR, or SAR file by adding it to a monitored directory is available only on distributed and z/OS operating systems. It is not supported on IBM i operating systems.

About this task

For applications previously installed by adding an EAR, JAR, WAR, or SAR file to a monitored *monitoredDeployableApps* subdirectory, you can uninstall the application or module from the application server by dragging or deleting the EAR, JAR, WAR, or SAR file from the *monitoredDeployableApps* subdirectory.

The product scans a monitored directory for new applications no more frequently than every five seconds, by default. The server or cluster must be running for the product to notice changes to files in its monitored directory.

After not finding the EAR, JAR, WAR, or SAR file in the monitored directory, the product stops the application or module and uninstalls the file from the application server or cluster. Also, the product removes application files from the directory for installed applications:

- *app_server_root/profiles/default/installedApps/cell_name*

Uninstalling an application file deletes it from the product configuration repository and deletes the application module binaries from the file system of the node where the application modules are installed.

This topic assumes that you use a graphical file browser to drag or delete the EAR, JAR, WAR, or SAR file from the directory. Alternatively, you can use operating system commands to delete a file from a monitored directory.

Procedure

1. Ensure that the application server or cluster from which you want to remove the enterprise application file is running.

To see which servers or clusters are running, you can use the `serverStatus -all` command. To start a server, you can use the `startServer server_name` command.

If the server is not running, start the server. For example, to start an application server named `server1` on `AppSrv02`, run the `startServer` command from a command prompt at the `app_server_root/profiles/AppSrv02/bin` directory:

```
startServer server1
```

The `Server server1 open for e-business message` indicates that the server is running.

2. Open a file browser and locate the monitored directory.
3. Remove the EAR, JAR, WAR, or SAR file that you want to uninstall.

Results

The product stops the application or module and uninstalls it from the application server or cluster.

Note: This topic references one or more of the application server log files. Beginning in WebSphere Application Server Version 8.0 you can configure the server to use the High Performance Extensible Logging (HPEL) log and trace infrastructure instead of using `SystemOut.log`, `SystemErr.log`, `trace.log`, and `activity.log` files or native z/OS logging facilities. If you are using HPEL, you can access all of your log and trace information using the LogViewer command-line tool from your server profile bin directory. See the information about using HPEL to troubleshoot applications for more information on using HPEL.

For uninstallations from a federated node of a deployment manager, if a node agent is stopped, the product uninstalls the application or module at the deployment manager level, but does not synchronize the changes with the node. The changes will occur at the node the next time the agent is running and a node synchronization occurs.

What to do next

Examine the server or cluster to verify that the application or module is no longer installed.

If the uninstallation is not successful, read messages in the `SystemOut.log` and fix the error condition. Add the EAR, JAR, WAR, or SAR file back into the monitored `monitoredDeployableApps` subdirectory, wait for the product to detect the file, and then remove it from the directory again. If the uninstallation still is not

successful, use the administrative console or a wsadmin script to uninstall the application file.

Deploying and administering applications: Resources for learning

Use the following links to find relevant supplemental information about deploying and administering 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”
- “Administration”

Programming model and decisions

- Designing Enterprise Applications with the Java™ 2 Platform, Enterprise Edition, Second Edition, http://java.sun.com/blueprints/guidelines/designing_enterprise_applications_2e/
- Java EE Tutorials, <http://java.sun.com/javaee/reference/tutorials/>
- Recommended reading list: J2EE and WebSphere Application Server, http://www.ibm.com/developerworks/websphere/library/techarticles/0305_issw/recommendedreading.html
- Java EE 5: Power and productivity with less complexity – An overview of Java EE 5 features and developer-productivity enhancements, <http://www.ibm.com/developerworks/java/library/j-jee5/index.html?ca=drs->
- Rational Application Developer V7 Programming Guide, SG24-7501-00, <http://www.redbooks.ibm.com/abstracts/sg247501.html?open>
- *IBM WebSphere Developer Technical Journal*: The top Java EE best practices, http://www.ibm.com/developerworks/websphere/techjournal/0701_botzum/0701_botzum.html

Programming instructions and examples

- *IBM WebSphere: Deployment and Advanced Configuration*, Roland Barcia, et al., ISBN 0131468626 (Prentice Hall, 2004)
- *IBM WebSphere Developer Technical Journal*: Co-hosting multiple versions of J2EE applications, http://www.ibm.com/developerworks/websphere/techjournal/0405_poddar/0405_poddar.html
- Automated Deployment of Enterprise Application Updates: Part 1 - Basic concepts, <http://websphere.sys-con.com/read/47889.htm>

Administration

- *IBM WebSphere Developer Technical Journal*: System management for WebSphere Application Server V6 -- Part 1 Overview of system management enhancements, http://www.ibm.com/developerworks/websphere/techjournal/0501_williamson/0501_williamson.html
- *IBM WebSphere Developer Technical Journal*: System management for WebSphere Application Server V6 -- Part 5: Flexible options for updating deployed applications, http://www.ibm.com/developerworks/websphere/techjournal/0510_apte/0510_apte.html
- *WebSphere Application Server V6.1: System Management Configuration Handbook*, SG24-7304-00, <http://www.redbooks.ibm.com/abstracts/SG247304.html?open>

Chapter 9. Managing applications through programming

Through Java MBean programming, you can install, update, and delete a Java Platform, Enterprise Edition (Java EE) application on a WebSphere Application Server deployment target.

Before you begin

This task assumes a basic familiarity with MBean programming. For information on MBean programming, see MBean Java application programming interface (API) documentation.

For information on the restarting of updated applications, refer to Fine-grained recycle behavior in *IBM WebSphere Developer Technical Journal: System management for WebSphere Application Server V6 -- Part 5 Flexible options for updating deployed applications*.

Before you can install or change an application on a deployment target, you must first create or update your application and assemble it using an assembly tool.

About this task

Besides installing, uninstalling, and updating applications through programming, you can additionally install, uninstall, and update Java EE applications through the administrative console or the wsadmin tool. All three ways provide identical updating capabilities.

Procedure

1. Perform any or all of the following tasks to manage your Java EE applications through programming.
 - Access the application management function.

This topic provides examples to access the application management functionality:

 - From WebSphere Application Server code
 - From outside WebSphere Application Server
 - When WebSphere Application Server is not running
 - Install an application.

This topic provides an example for initially installing an application on a deployment target such as a server or cluster.
 - Uninstall an application.

This topic provides an example for uninstalling an application that resides on a deployment target.
 - Manipulate additional attributes for a deployed application.

This topic provides an example for manipulating attributes that are not exposed through the `AppDeploymentTask` object.
 - Share sessions for application management.

This topic provides an example for saving application-specific updates for a deployed application to a session, and then to the configuration repository.
 - Update an application.

This topic provides an example for updating the installed application on a server or cluster with a new application. When you completely update an application, the deployed application is uninstalled and the new enterprise archive (EAR) file is installed.
 - Add to, update, or delete part of an application.

This topic provides an example that you can use to add, update, or delete part of an application on a server or cluster.
 - Edit an application.

This topic provides an example that you can use to edit an application on a server or cluster.

- Add a module.
This topic provides an example for adding a module to an application that resides on a server or cluster.
- Update a module.
This topic provides an example for updating a module that resides on a server or cluster. When you update a module, the deployed module is uninstalled and the updated module is installed.
- Delete a module.
This topic provides an example for deleting a module that resides on a server or cluster. When you delete a module, the deployed module is uninstalled.
- Add a file.
This topic provides an example for adding a file to an application that resides on a server or cluster.
- Update a file.
This topic provides an example for updating a file on a server or cluster. When you update a file, the deployed file is uninstalled and the updated file is installed.
- Delete a file.
This topic provides an example for deleting a file on a server or cluster. When you delete a file, the deployed file is uninstalled.

2. Save your changes to the master configuration repository.
3. Synchronize changes to the master configuration across the nodes for the changes to take effect.

What to do next

If you have further application updates, you can do the updates through programming, the administrative console, or the wsadmin tool.

You can use the common deployment framework to add additional logic to application management operations. See Chapter 10, “Extending application management operations through programming,” on page 383. The tasks that the extensions provide are available through all the administrative clients, such as the wsadmin tool, the administrative console, or through programmatic APIs that the AppManagement MBean provides.

Accessing the application management function

The `com.ibm.websphere.management.application.AppManagementProxy` class provides uniform access to application management functionality, regardless of whether the functionality is accessed from the server process, administrative client process, or a stand-alone Java program in the absence of WebSphere Application Server. This topic provides code excerpts that demonstrate how to obtain an `AppManagementProxy` instance in a variety of cases.

Before you begin

This task assumes a basic familiarity with WebSphere Application Server programming interfaces and MBean programming. Read about WebSphere Application Server programming interfaces and MBean programming in the application programming interfaces documentation.

About this task

Perform any of the following tasks to access application management functionality through programming.

Procedure

- To access application management functionality from WebSphere Application Server code, for example, as a custom service, create the `AppManagementProxy` class.

```
AppManagement appMgmt =
    AppManagementProxy.getJMXProxyForServer();
```

- To access application management functionality from outside WebSphere Application Server through the AppManagement MBean, create an administrative client to establish a connection to WebSphere Application Server and then create the AppManagementProxy class.

```
AdminClient adminClient = ....
```

```
// create AppManagement proxy object
AppManagement appMgmt = AppManagementProxy.getJMXProxyForClient (adminClient);
```

- To access application management functionality when WebSphere Application Server is not running (local mode), create the AppManagementProxy class.

```
AppManagement appMgmt = AppManagementProxy.getLocalProxy ();
```

- When running in local mode set the `com.ibm.ws.management.standalone` system property to `true`. If you want to modify configuration documents in a non-default location, set the location of the configuration directory through the `was.repository.root` system property.
- Although you can use application management functions with or without WebSphere Application Server running, do not access application management functions concurrently through local mode and the AppManagement MBean. Otherwise, updates that are made using these modes can collide and break the integrity of the WebSphere Application Server configuration.

Results

After you successfully create the AppManagementProxy class, you have access to application management functionality.

What to do next

You can perform various management tasks such as installing, uninstalling, editing, and so on.

Installing an application through programming

You can install an application through the administrative console, the wsadmin tool, or programming. Use this example to install an application through programming.

Before you begin

This task assumes a basic familiarity with MBean programming. For information on MBean programming, see MBean Java application programming interface (API) documentation.

Before you can install an application on WebSphere Application Server, you must install the application.

About this task

The AppDeploymentController instance contains meta-data defined in XMLI-based deployment descriptors as well as annotations defined in Java classes within the input enterprise archive (EAR) file.

Perform the following tasks to install an application through programming.

Procedure

1. Populate the EAR file with WebSphere Application Server-specific binding information.
 - a. Create the controller and populate the EAR file with appropriate options.
 - b. Optionally run the default binding generator.
 - c. Save and close the EAR file.
 - d. Retrieve the saved options table that will be passed to the installApplication MBean API.
2. Connect to WebSphere Application Server.

3. Create the application management proxy.
4. If the preparation phase (population of the EAR file) is not performed, then do the following actions:
 - a. Create an options table to be passed to the installApplication MBean API.
 - b. Create a table for module to server relations and add the table to the options table.
Refer to the `com.ibm.websphere.management.application.AppManagement` class in the application programming interfaces documentation to understand various options that can be passed to the `installApplication` MBean API.
5. Create the notification filter for listening to installation events.
6. Add the listener.
7. Install the application.
8. Wait for some timeout so that the program does not end.
9. Listen to Java Management Extensions (JMX) notifications to understand completion of the operation.
10. When the installation is done, remove the listener and quit.

Results

After you successfully run the code, the application is installed.

Example

The following example shows how to install an application based on the previous steps. Some statements are split on multiple lines for printing purposes.

```
import java.lang.*;
import java.io.*;
import java.util.*;
import java.lang.reflect.*;
import com.ibm.websphere.management.application.*;
import com.ibm.websphere.management.application.client.*;
import com.ibm.websphere.management.*;

import javax.management.*;

public class Install {

    public static void main (String [] args) {

        try {
            String earFile = "C:/test/test.ear";
            String appName = "MyApp";

            // Preparation phase: Begin
            // Through the preparation phase you populate the enterprise archive (EAR) file with
            // WebSphere Application Server-specific binding information. For example, you can specify
            // Java Naming and Directory Interface (JNDI) names for enterprise beans, or virtual hosts
            // for web modules, and so on.

            // First, create the controller and populate the EAR file with the appropriate options.
            Hashtable prefs = new Hashtable();
            prefs.put(AppConstants.APPDEPL_LOCALE, Locale.getDefault());

            // You can optionally run the default binding generator by using the following options.
            // Refer to Java documentation for the AppDeploymentController class to see all the
            // options that you can set.
            Properties defaultBnd = new Properties();
            prefs.put (AppConstants.APPDEPL_DFLTBNDRG, defaultBnd);
            defaultBnd.put (AppConstants.APPDEPL_DFLTBNDRG_VHOST, "default_host");

            // Create the controller.
            AppDeploymentController controller = AppDeploymentController
                .readArchive(earFile, prefs);
            AppDeploymentTask task = controller.getFirstTask();
            while (task != null)
            {
                // Populate the task data.
                String[][] data = task.getTaskData();
                // Manipulate task data which is a table of stringtask.
                task.setTaskData(data);
                task = controller.getNextTask();
            }
            controller.saveAndClose();

            Hashtable options = controller.getAppDeploymentSavedResults();
```

```

// The previous options table contains the module-to-server relationship if it was set by
// using tasks.
//Preparation phase: End

// Get a connection to WebSphere Application Server.
String host = "localhost";
String port = "8880";
String target = "WebSphere:cell=cellName,node=nodeName,server=server1";

Properties config = new Properties();
config.put (AdminClient.CONNECTOR_HOST, host);
config.put (AdminClient.CONNECTOR_PORT, port);
config.put (AdminClient.CONNECTOR_TYPE, AdminClient.CONNECTOR_TYPE_SOAP);
System.out.println ("Config: " + config);
AdminClient _soapClient = AdminClientFactory.createAdminClient(config);

// Create the application management proxy, AppManagement.
AppManagement proxy = AppManagementProxy.getJMXProxyForClient (_soapClient);

// If code for the preparation phase has been run, then you already have the options table.
// If not, create a new table and add the module-to-server relationship to it by uncommenting
// the next statement.
//Hashtable options = new Hashtable();
options.put (AppConstants.APPDEPL_LOCALE, Locale.getDefault());

// Uncomment the following statements to add the module to the server relationship table if
// the preparation phase does not collect it.
//Hashtable module2server = new Hashtable();
//module2server.put ("*", target);
//options.put (AppConstants.APPDEPL_MODULE_TO_SERVER, module2server);

//Create the notification filter for listening to installation events.
NotificationFilterSupport myFilter = new NotificationFilterSupport();
myFilter.enableType (AppConstants.NotificationType);

//Add the listener.
NotificationListener listener = new AListener(_soapClient,
myFilter, "Install: " + appName, AppNotification.INSTALL);

// Install the application.
proxy.installApplication (earFile, appName, options, null);
System.out.println ("After install App is called..");

// Wait for some timeout. The installation application programming interface (API) is
// asynchronous and so returns immediately.
// If the program does not wait here, the program ends.
Thread.sleep(300000); // Wait so that the program does not end.

    }
    catch (Exception e) {
        e.printStackTrace();
    }
}

}

// Specify the Java Management Extensions (JMX) notification listener for JMX events.
class AListener implements NotificationListener
{
    AdminClient _soapClient;
    NotificationFilterSupport myFilter;
    Object handback;
    ObjectName on;
    String eventTypeToCheck;

    public AListener(AdminClient cl, NotificationFilterSupport fl,
Object h, String eType) throws Exception
    {
        _soapClient = cl;
        myFilter = fl;
        handback = h;
        eventTypeToCheck = eType;

        Iterator iter = _soapClient.queryNames (new ObjectName(
"WebSphere:type=AppManagement,*"), null).iterator();
        on = (ObjectName)iter.next();
        System.out.println ("ObjectName: " + on);
        _soapClient.addNotificationListener (on, this, myFilter, handback);
    }

    public void handleNotification (Notification notf, Object handback)
    {
        AppNotification ev = (AppNotification) notf.getUserData();
        System.out.println ("!! JMX event Recd: (handback obj= " + handback+ "): " + ev);

        //When the installation is done, remove the listener and quit.

```

```

    if (ev.taskName.equals (eventTypeToCheck) &&
        (ev.taskStatus.equals (AppNotification.STATUS_COMPLETED) ||
         ev.taskStatus.equals (AppNotification.STATUS_FAILED)))
    {
        try
        {
            _soapClient.removeNotificationListener (on, this);
        }
        catch (Throwable th)
        {
            System.out.println ("Error removing listener: " + th);
        }
        System.exit (0);
    }
}
}
}

```

What to do next

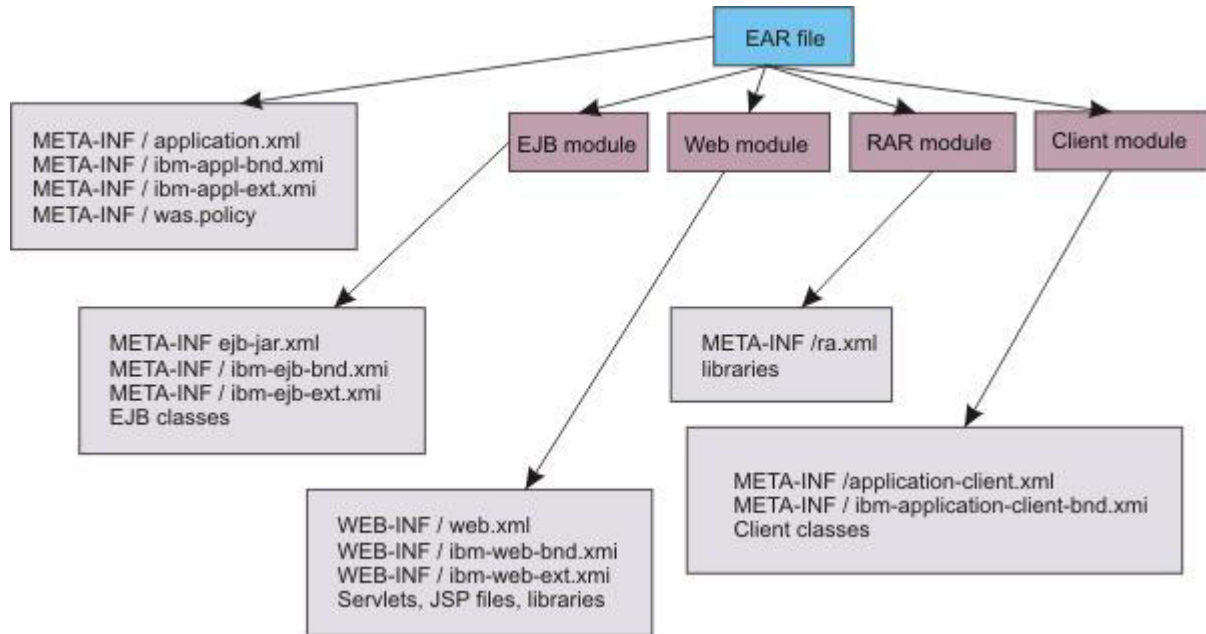
Once you install the application, you must explicitly start the application or you must stop and restart the server. For information on starting an application, see the Starting an application through programming topic in the *Using the administrative clients* PDF. For information on stopping or restarting the server, see the Stopping an application server topic or the Starting an application server topic, respectively, in the *Setting up the application serving environment* PDF.

Application management

Java Platform, Enterprise Edition (Java EE) applications and modules include an Extensible Markup Language (XML)-based deployment descriptor that specifies various Java EE artifacts that pertain to applications or modules. The Java EE artifacts include Enterprise JavaBeans (EJB) definitions, security role definitions, EJB references, resource references, and so on. These artifacts define various unresolved references that the application logic uses. The Java EE specification requires that these artifacts map to Java EE platform-specific information, such as that found in WebSphere Application Server, during deployment of Java EE applications.

The application assembly tools that WebSphere Application Server supports, as well as the application management support that is provided with the product, facilitate collection of certain WebSphere Application Server information. The collected information is used to resolve references that are defined in various deployment descriptors in a Java EE application. This information is stored in the application EAR file in conjunction with the deployment descriptors. The following diagram shows the structure of an enterprise archive (EAR) file that is populated with deployment information that is specific to WebSphere

Application Server.



The application management architecture provides a set of classes with which deployers can collect WebSphere Application Server deployment information. This information is also referred to as *binding information*, and is stored in the application EAR file. The deployer can install the EAR file into a WebSphere Application Server configuration by using the AppManagement interface.

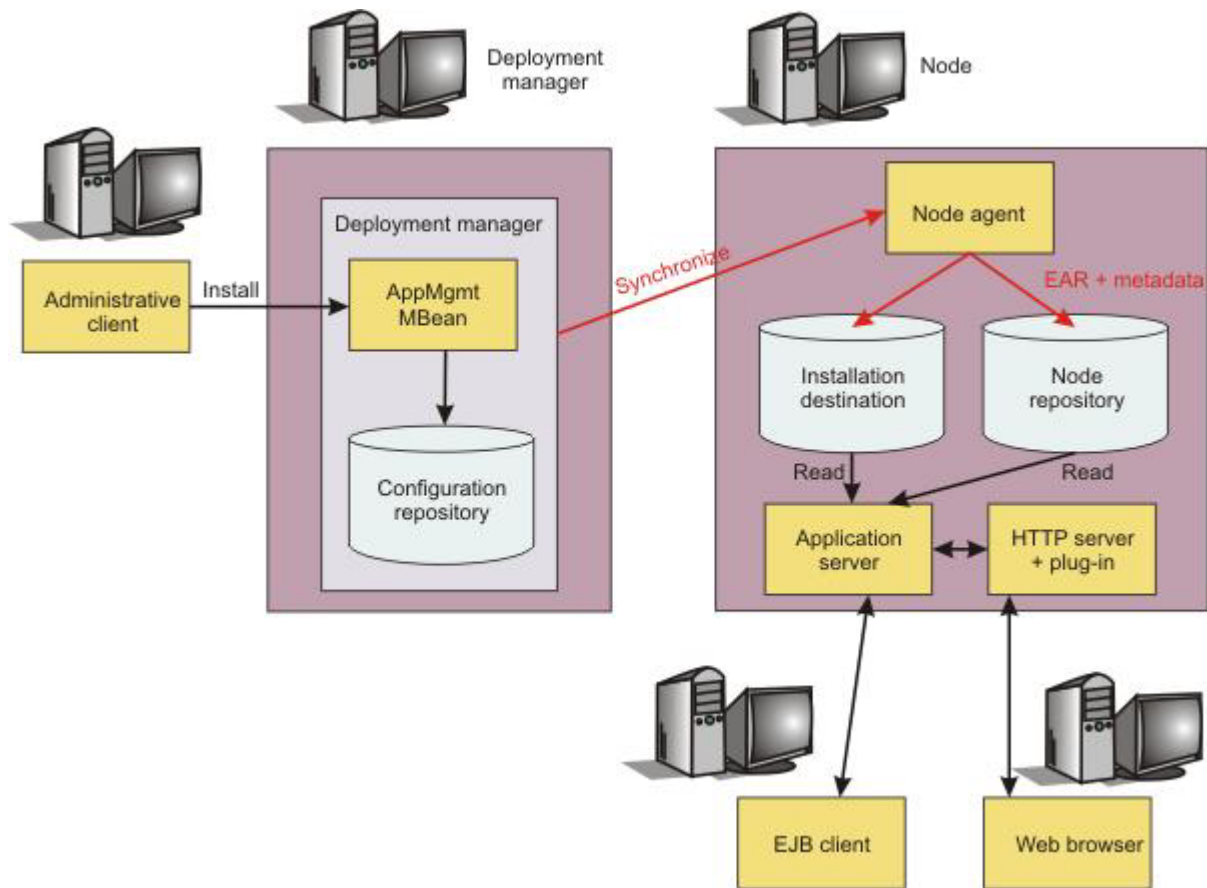
The application management support in WebSphere Application Server provides functions such as installing and uninstalling applications, editing binding information for installed applications, updating the entire application or part of the application, exporting the application, and so on. The `com.ibm.websphere.management.application.AppManagement` interface, which is exposed as a Java Management Extensions (JMX)-based AppManagement MBean in WebSphere Application Server, provides this functionality. Code that runs on the server or in a stand-alone administrative client program can access the interface. Access to the application management functions is also possible in the absence of WebSphere Application Server. This mode, known as *local mode*, is particularly useful for installing Java EE applications as part of product installation. For WebSphere Application Server, Network Deployment, the AppManagement MBean is present in the deployment manager only, which facilitates centralized configuration and administration.

Application deployment in a WebSphere Application Server, Network Deployment configuration

The following diagram explains application deployment for the WebSphere Application Server, Network Deployment product:

1. The application EAR file that you install is stored in the master configuration repository of the deployment manager, with the application metadata.
2. The configuration synchronization operation replicates the application in the master repository into the repository of the target node.
3. The application EAR file is extracted to the installation destination of the target node at the end of the synchronization operation.
4. The WebSphere Application Server run time reads the application business logic such as EJB classes, servlets, and JavaServer Pages (JSP) from the installation destination while serving application client requests.

- WebSphere Application Server reads the application metadata, such as deployment descriptors and WebSphere Application Server bindings during application startup from the configuration repository or the installation destination only, depending on the option that is specified during application installation.



Starting an application through programming

You can start an application through the administrative console, the wsadmin tool, or programming. Use this example to start an application through programming.

Before you begin

This task assumes a basic familiarity with MBean programming. For information on MBean programming, see MBean Java application programming interface (API) documentation.

Before you can start an application on WebSphere Application Server, you must install your application.

About this task

Perform the following tasks to start an application through programming.

Procedure

- Connect the administrative client to WebSphere Application Server.
- Create the application management proxy.
- Call the startApplication method on the proxy by passing the application name and optionally the list of targets on which to start the application.

Results

After you successfully run the code, the application is started.

Example

The following example shows how to start an application following the previously listed steps. Some statements are split on multiple lines for printing purposes.

```
//Do a get of the administrative client to connect to
//WebSphere Application Server.

AdminClient client = ...;
String appName = "myApp";
Hashtable prefs = new Hashtable();
// Use the AppManagement MBean to start and stop applications on all or some targets.
// The AppManagement MBean is on the deployment manager in the WebSphere Application Server, Network Deployment product.

// Query and get the AppManagement MBean.
ObjectName on = new ObjectName ("WebSphere:type=AppManagement,*");
Iterator iter = client.queryNames (on, null).iterator();
ObjectName appmgmtON = (ObjectName)iter.next();

//Start the application on all targets.
AppManagement proxy = AppManagementProxy.getJMXProxyForClient(client);
String started = proxy.startApplication(appName, prefs, null);
System.out.println("Application started on folloing servers: " + started);

//Start the application on some targets.
//String targets = "WebSphere:cell=cellname,node=nodename,
server=servername+WebSphere:cell=cellname,cluster=clusterName";
//String started1 = proxy.startApplication(appName, targets, prefs, null);
//System.out.println("Application started on following servers: " + started1)
```

Sharing sessions for application management

With the configuration service interface, `ConfigService`, you can create a session that is a temporary staging area, where you can save all the configuration modifications. Saving the session saves all the updates from the session into the WebSphere Application Server configuration repository. The application management logic supports session sharing with the configuration service. You can perform all the application management functions in the same session as the one that the configuration service creates. Saving such a session saves all the updates, including the ones that are application-specific.

Before you begin

This task assumes a basic familiarity with MBean programming. For information on MBean programming, see [MBean Java application programming interface \(API\) documentation](#).

About this task

Perform the following tasks for your deployed application to share and save application-specific updates through the `configService` configuration service.

Procedure

1. Create a configuration service proxy object.
2. Create a session.
3. Pass the session information to the `AppManagement` MBean.

Every method on the `AppManagement` interface takes session ID (workspace ID) as the last parameter. If the session information is passed in this parameter, the application management function uses the session. If you set the parameter to a null value:

- No session sharing occurs
- The configuration changes are always saved in the configuration repository if the operation succeeds.

4. Save the session after all the necessary changes are made.

The following example outlines the general steps for session sharing through the configService configuration service. For a detailed example, see “Manipulating additional attributes for a deployed application.”

```
public void installApplication (String localEarPath,
                               String appName, Hashtable properties, String workspaceID)
    throws AdminException;

AdminClient adminClient = ....;

// Create a configuration service proxy object.
ConfigService configService = new ConfigServiceProxy(adminClient);

// Create a session.
Session session = new Session();

// Pass the session information to AppManagement MBean.
appMgmt = ...
appMgmt.installApplication
    (earPath, appName, properties, session.toString());
//Save the session after all necessary changes are made.
configService.save(session, false);
```

Results

After you successfully complete the steps, you have saved application-specific updates for a deployed application to a session, and then to the configuration repository.

Manipulating additional attributes for a deployed application

You can manipulate attributes for a deployed application through the administrative console, the wsadmin tool, or by programming. Use this example to manipulate attributes that are not exposed during or after application installation through the AppDeploymentTask object.

Before you begin

This task assumes a basic familiarity with MBean programming and the ConfigService interfaces. Read about MBean programming and ConfigService interfaces in the application programming interfaces documentation.

About this task

Perform the following tasks for your deployed application to manipulate attributes that are not exposed through the AppDeploymentTask object. The attributes are saved in the deployment.xml file that is created in the configuration repository for each deployed application.

Procedure

1. Create a session.
2. Connect to WebSphere Application Server.
3. Locate the ApplicationDeployment object.
4. Manipulate the attributes.
5. Save your changes.
6. Clean up the session.

Results

After you successfully run the code, the attributes are updated in the `deployment.xml` file for the deployed application.

Example

The following example shows how to manipulate the `startingWeight`, `warClassLoaderPolicy`, and `classloader` attributes based on the previous steps.

```
import java.util.Properties;

import javax.management.Attribute;
import javax.management.AttributeList;
import javax.management.ObjectName;

import com.ibm.websphere.management.AdminClient;
import com.ibm.websphere.management.AdminClientFactory;
import com.ibm.websphere.management.Session;
import com.ibm.websphere.management.configservice.ConfigService;
import com.ibm.websphere.management.configservice.ConfigServiceHelper;
import com.ibm.websphere.management.configservice.ConfigServiceProxy;
import com.ibm.websphere.management.exception.ConfigServiceException;
import com.ibm.websphere.management.exception.ConnectorException;

public class AppManagementSample1 {

    public static void main(String[] args) {
        String hostName = "localhost";
        String port = "8880";
        String appName = "ivtApp";

        ConfigService configService = null;

        // create a session.
        Session session = new Session();

        // establish connection to the server.
        try {
            Properties props = new Properties();
            props.setProperty(AdminClient.CONNECTOR_TYPE,
                AdminClient.CONNECTOR_TYPE_SOAP);
            props.setProperty(AdminClient.CONNECTOR_HOST, hostName);
            props.setProperty(AdminClient.CONNECTOR_PORT, port);
            AdminClient adminClient =
                AdminClientFactory.createAdminClient(props);

            // create a config service proxy object.
            configService = new ConfigServiceProxy(adminClient);

            // Locate the application object.
            ObjectName rootID = configService.resolve(session,
                "Deployment="+appName)[0];
            System.out.println ("rootID is: " + rootID);

            // Locate the ApplicationDeployment object from the root.
            ObjectName appDeplPattern = ConfigServiceHelper.createObjectName
                (null, "ApplicationDeployment");
            /*
            ObjectName appDeplID = configService.queryConfigObjects(session,
                rootID, appDeplPattern, null)[0];
            */
            AttributeList list1 = configService.getAttributes(session,
                rootID, new String[]{"deployedObject"}, false);
            ObjectName appDeplID = (ObjectName)
                ConfigServiceHelper.getAttributeValue(list1, "deployedObject");
            System.out.println ("appDeplID: " + appDeplID);

            // Locate the class loader.

            // Change the starting weight through the startingWeight attribute. The starting weight
            // affects the order in which applications start.
            AttributeList attrList = new AttributeList();
            Integer newWeight = new Integer (10);
            attrList.add(new Attribute("startingWeight", newWeight));

            // Change the WAR class loader policy through the warClassLoaderPolicy attribute by
            // specifying SINGLE or MULTIPLE.
            // SINGLE=one classloader for all WAR modules
            attrList.add(new Attribute("warClassLoaderPolicy", "SINGLE"));

            // Set the class loader mode to PARENT_FIRST or PARENT_LAST.
            AttributeList clList = (AttributeList) configService.getAttribute
                (session, appDeplID, "classloader");
            ConfigServiceHelper.setAttributeValue (clList, "mode",
```

```

    "PARENT_LAST");
    attrList.add (new Attribute ("classloader", clList));

    // Set the new values.
    configService.setAttributes(session, appDepID, attrList);

    // Save your changes.
    configService.save(session, false);

} catch (Exception ex) {
    ex.printStackTrace();
} finally {
    // Clean up the session.
    try {
        configService.discard(session);
    }
    catch (ConfigServiceException csEx)
    {
        csEx.printStackTrace();
    }
    catch (ConnectorException cnEx)
    {
        cnEx.printStackTrace();
    }
}
}
}

```

Editing applications

You can edit deployed applications through the administrative console, the wsadmin tool, or by programming. Use this example to edit a deployed application through programming.

Before you begin

This task assumes a basic familiarity with MBean programming. For information on MBean programming, see MBean Java application programming interface (API) documentation.

Before you can edit an application on WebSphere Application Server, you must install the application.

About this task

Perform the following tasks to edit your deployed application.

Procedure

1. Connect to WebSphere Application Server.
2. Create the application management proxy.
3. Get information about an installed application.
4. Manipulate task data as necessary.
5. Save changes for the installed application.

Results

After you successfully run the code, the application is edited.

Example

The following example shows how to edit an application, based on the previous steps.

```

import java.lang.*;
import java.io.*;
import java.util.*;
import java.lang.reflect.*;
import com.ibm.websphere.management.application.*;
import com.ibm.websphere.management.application.client.*;
import com.ibm.websphere.management.*;

import javax.management.*;

```

```

public class aa {

    public static void main (String [] args) {

        try {

            // Connect to WebSphere Application Server.
            String host = "localhost";
            String port = "8888";
            String target = "WebSphere:cell=cellName,node=nodeName,server=server1";

            Properties config = new Properties();
            config.put (AdminClient.CONNECTOR_HOST, host);
            config.put (AdminClient.CONNECTOR_PORT, port);
            config.put (AdminClient.CONNECTOR_TYPE, AdminClient.CONNECTOR_TYPE_SOAP);
            System.out.println ("Config: " + config);
            AdminClient _soapClient = AdminClientFactory.createAdminClient(config);

            // Create the application management proxy, AppManagement.
            AppManagement proxy = AppManagementProxy. getJMXProxyForClient (_soapClient);

            String appName = "MyApp";
            // Get information for an application with name appName:
            // Pass Locale information as the preference.
            Hashtable prefs = new Hashtable();
            prefs.put(AppConstants.APPDEPL_LOCALE, Locale.getDefault());
            Vector allTasks = appMgmt.getApplicationInfo (appName, prefs, null);

            // Manipulate task data as necessary.
            if (task.getName().equals ("MapRolesToUsers") && !task. isTaskDisabled())
            {

                // find out column index for role and user column
                // refer to the previous table to find the column names
                int roleColumn = -1;
                int userColumn = -1;
                String[] colNames = task.getColumnNames();
                for (int i=0; i < colNames.length; i++)
                {
                    if (colNames[i].equals ("role")
                        roleColumn = i;
                    else if (colNames[i].equals ("role.user"))
                        userColumn = i;
                }

                // iterate over task data starting at row 1 as row0 is
                // column names
                String[][] data = task.getTaskData();
                for (int i=1; i < data.length; i++)
                {
                    if (data[i][roleColumn].equals ("Role1"))
                    {
                        data[i][userColumn]="User1|User2";
                        break;
                    }
                }

                // now that the task data is changed, save it back
                task.setTaskData (data);
            }

            // Save changes back into the installed application:
            // Set information for an application with name appName.
            // Pass Locale information as the preference.
            prefs = new Hashtable();
            prefs.put(AppConstants.APPDEPL_LOCALE, Locale.getDefault());
            appMgmt.setApplicationInfo (appName, prefs, null, allTasks);

            }
        catch (Exception e) {
            e.printStackTrace();
        }
    }
}

```

Updating an application through programming

You can update an existing application through the administrative console, the wsadmin tool, or programming. Use this example to completely update an application through programming.

Before you begin

This task assumes a basic familiarity with MBean programming. For information on MBean programming, see MBean Java application programming interface (API) documentation.

Before you can update an application on WebSphere Application Server, you must install your application.

About this task

Perform the following tasks to completely update an application through programming.

Procedure

1. Connect to WebSphere Application Server.
2. Create the application management proxy.
3. Create the notification filter for listening to events.
4. Add the listener.
5. Prepare the enterprise archive (EAR) file by populating it with binding information.
6. Update the application.
7. Wait for some timeout so that the program does not end.
8. Listen to Java Management Extensions (JMX) notifications to understand completion of the operation.
9. After you update the application, remove the listener and quit.

Results

After you successfully run the code, the application is updated.

Example

The following example shows how to update an application based on the previous steps. Some statements are split on multiple lines for printing purposes.

```
import java.lang.*;
import java.io.*;
import java.util.*;
import java.lang.reflect.*;
import com.ibm.websphere.management.application.*;
import com.ibm.websphere.management.application.client.*;
import com.ibm.websphere.management.*;

import javax.management.*;

public class aa {

    public static void main (String [] args) {

        try {

            // Connect to WebSphere Application Server.
            String host = "localhost";
            String port = "8880";
            String target = "WebSphere:cell=cellName,node=nodeName,server=server1";

            Properties config = new Properties();
            config.put (AdminClient.CONNECTOR_HOST, host);
            config.put (AdminClient.CONNECTOR_PORT, port);
            config.put (AdminClient.CONNECTOR_TYPE, AdminClient.CONNECTOR_TYPE_SOAP);
            System.out.println ("Config: " + config);
            AdminClient _soapClient = AdminClientFactory.createAdminClient(config);

            // Create the application management proxy, AppManagement.
            AppManagement proxy = AppManagementProxy.getJMXProxyForClient (_soapClient);

            String appName = "MyApp";
            String fileContents = "/test/test.ear";

            // Create the notification filter.
            NotificationFilterSupport myFilter = new NotificationFilterSupport();
            myFilter.enableType (NotificationConstants.TYPE_APPMANAGEMENT);
            //Add the listener.
```



```

NotificationListener listener = new AListener(_soapClient, myFilter,
"Install: " + appName, AppNotification.INSTALL);

// Refer to the installation example to see how you can prepare the enterprise archive (EAR)
// file by populating it with binding information.
// If code for the preparation phase has started, then you already have the options table.
// If not, create a new table and add the module-to-server relationship to it by uncommenting
// the next statement.
//Hashtable options = new Hashtable();
options.put (AppConstants.APPDEPL_LOCALE, Locale.getDefault());
options.put ((AppConstants.APPUPDATE_CONTENTTYPE, AppConstants.APPUPDATE_CONTENT_APP);

// Uncomment the following statements to add the module to the server relationship table if
// the preparation phase does not collect it
//Hashtable module2server = new Hashtable();
//module2server.put ("*", target);
//options.put (AppConstants.APPDEPL_MODULE_TO_SERVER, module2server);
// Update the application.
proxy.updateApplication (appName,
null,
fileContents,
AppConstants.APPUPDATE_UPDATE,
options,
null);

// Wait for some timeout. The installation application programming interface (API) is
// asynchronous and so returns immediately.
// If the program does not wait here, the program ends.
Thread.sleep(300000); // Wait so that the program does not end.

    }
    catch (Exception e) {
        e.printStackTrace();
    }
}
}

// Specify the Java Management Extensions (JMX) notification listener for JMX events.
class AListener implements NotificationListener
{
    AdminClient _soapClient;
    NotificationFilterSupport myFilter;
    Object handback;
    ObjectName on;
    String eventTypeToCheck;

    public AListener(AdminClient cl, NotificationFilterSupport fl,
Object h, String eType) throws Exception
    {
        _soapClient = cl;
        myFilter = fl;
        handback = h;
        eventTypeToCheck = eType;

        Iterator iter = _soapClient.queryNames (new ObjectName(
"WebSphere:type=AppManagement,*"), null).iterator();
on = (ObjectName)iter.next();
System.out.println ("ObjectName: " + on);
_soapClient.addNotificationListener (on, this, myFilter, handback);
    }

    public void handleNotification (Notification notf, Object handback)
    {
        AppNotification ev = (AppNotification) notf.getUserData();
        System.out.println ("!! JMX event Recd: (handback obj= " + handback+ "): " + ev);

        //When the installation is done, remove the listener and quit

        if (ev.taskName.equals (eventTypeToCheck) &&
(ev.taskStatus.equals (AppNotification.STATUS_COMPLETED) ||
ev.taskStatus.equals (AppNotification.STATUS_FAILED)))
        {
            try
            {
                _soapClient.removeNotificationListener (on, this);
            }
            catch (Throwable th)
            {
                System.out.println ("Error removing listener: " + th);
            }
            System.exit (0);
        }
    }
}
}

```

Adding to, updating, or deleting part of an application through programming

You can add to, update, or delete part of an existing application through the administrative console, the wsadmin tool, or programming. This example changes part of an application through programming. You can use this example whether you add to, update, or delete part of an existing application. Multiple changes to an application can be packaged in a single compressed .zip file.

Before you begin

To learn about the structure of the compressed .zip file, see the Updating applications topic in the *Developing and deploying applications* PDF. This task assumes a basic familiarity with MBean programming. For information on MBean programming, see MBean Java application programming interface (API) documentation.

Before you can add to, update, or delete part of an application on WebSphere Application Server, you must install the application.

About this task

Perform the following tasks to add to, update, or delete part of an application through programming.

Procedure

1. Connect to WebSphere Application Server.
2. Create the application management proxy.
3. Create the notification filter.
4. Add the listener.
5. Partially change the existing application.
6. Wait for some timeout so that the program does not end.
7. Listen to Java Management Extensions (JMX) notifications to understand completion of the operation.

Results

After you successfully run the code, you have changed the application.

Example

The following example shows how to add to, update, or delete part of an application based on the previous steps. Some statements are split on multiple lines for printing purposes.

```
//Inputs:
//partialApp specifies the location of the partial application.
//appName specifies the name of the application.
String partialApp = "/apps/partial.zip";
String appName = "MyApp";

//Do a get of the administrative client to connect to the product.
AdminClient client = ...;

//Create the application management proxy.
AppManagement proxy = AppManagementProxy.getJMXProxyForClient (client);

// Create the notification filter.
NotificationFilterSupport myFilter = new NotificationFilterSupport();
myFilter.enableType (NotificationConstants.TYPE_APPMANAGEMENT);
//Add the listener.
NotificationListener listener = new AListener(_soapClient, myFilter,
"Install: " + appName, AppNotification.UPDATE);
//Partially change the existing application, MyApp.

Hashtable options = new Hashtable();
```

```

options.put (AppConstants.APPDEPL_LOCALE, Locale.getDefault());
options.put (AppConstants.APPUPDATE_CONTENTTYPE, AppConstants.APPUPDATE_CONTENT_PARTIALAPP);

proxy.updateApplication ( appName,
    null,
    partialApp,
    null,
    options,
    null);

// Wait for some timeout. The installation application programming interface (API) is
// asynchronous and so returns immediately.
// If the program does not wait here, the program ends.
Thread.sleep(300000); // Wait so that the program does not end.
}
catch (Exception e) {
    e.printStackTrace();
}
}
}
// Specify the Java Management Extensions (JMX) notification listener for JMX events.
class AListener implements NotificationListener
{
    AdminClient _soapClient;
    NotificationFilterSupport myFilter;
    Object handback;
    ObjectName on;
    String eventTypeToCheck;

    public AListener(AdminClient cl, NotificationFilterSupport fl,
Object h, String eType) throws Exception
    {
        _soapClient = cl;
        myFilter = fl;
        handback = h;
        eventTypeToCheck = eType;

        Iterator iter = _soapClient.queryNames (new ObjectName(
"WebSphere:type=AppManagement,*"), null).iterator();
        on = (ObjectName)iter.next();
        System.out.println ("ObjectName: " + on);
        _soapClient.addNotificationListener (on, this, myFilter, handback);
    }

    public void handleNotification (Notification notf, Object handback)
    {
        AppNotification ev = (AppNotification) notf.getUserData();
        System.out.println ("!! JMX event Recd: (handback obj= " + handback+ "): " + ev);

        //When the installation is done, remove the listener and quit

        if (ev.taskName.equals (eventTypeToCheck) &&
            (ev.taskStatus.equals (AppNotification.STATUS_COMPLETED) ||
            ev.taskStatus.equals (AppNotification.STATUS_FAILED)))
        {
            try
            {
                _soapClient.removeNotificationListener (on, this);
            }
            catch (Throwable th)
            {
                System.out.println ("Error removing listener: " + th);
            }
            System.exit (0);
        }
    }
}
}

```

What to do next

After you update the application, remove the listener and quit.

Preparing a module and adding it to an existing application through programming

You can add a module to an existing application through the administrative console, the wsadmin tool, or programming. Use this example to add a module through programming.

Before you begin

This task assumes a basic familiarity with MBean programming. For information on MBean programming, see MBean Java application programming interface (API) documentation.

Before you can add a module to an application on WebSphere Application Server, you must install your application.

About this task

Perform the following tasks to add a module to an application through programming.

Procedure

1. Create an application deployment controller instance to populate the module file with binding information.
2. Save the binding information in the module.
3. Get the installation options.
4. If the preparation phase (population of the EAR file) is not performed, do the following actions:
 - a. Create an options table to be passed to the updateApplication MBean API.
 - b. Create a table for module to server relations and add the table to the options table.
5. Connect to WebSphere Application Server.
6. Create the application management proxy.
7. Create the notification filter.
8. Add the listener.
9. Add the module to the application.
10. Specify the target for the new module.
11. Wait for some timeout so that the program does not end.
12. Listen to Java Management Extensions (JMX) notifications to understand completion of the operation.

Results

After you successfully run the code, the module is added to the application.

Example

The following example shows how to add a module to an application based on the previous steps. Some statements are split on multiple lines for printing purposes.

```
//Inputs:
//moduleName specifies the name of the module that you add to the application.
//moduleURI specifies a URI that gives the target location of the module
// archive contents on a file system. The URI provides the location of the new
// module after installation. The URI is relative to the application URL.
//uniquemoduleURI specifies the URI that gives the target location of the
// deployment descriptor file. The URI is relative to the application URL.
//target specifies the cell, node, and server on which the module is installed.
String moduleName = "/apps/foo.jar";
String moduleURI = "Increment.jar";
String uniquemoduleURI = "Increment.jar+META-INF/ejb-jar.xml";
String target = "WebSphere:cell=cellname,node=nodename,server=servername";

//Create an application deployment controller instance, AppDeploymentController,
//to populate the Java Archive (JAR) file with binding information.
//The binding information is WebSphere Application Server-specific deployment information.

Hashtable preferences = new Hashtable();
preferences.put (AppConstants.APPDEPL_LOCALE, Locale.getDefault());
preferences.put (AppConstants.APPUPDATE_CONTENTTYPE, AppConstants.APPUPDATE_CONTENT_MODULEFILE);
AppDeploymentController controller = AppManagementFactory.readArchiveForUpdate(
    moduleName,
```

```

moduleURI,
AppConstants.APPUPDATE_ADD,
preferences,
null);

```

If the module that you add to the application lacks any bindings, add the bindings so that the module addition works. Collect and add the bindings by using the public APIs provided with the product. Refer to Java documentation for the `com.ibm.websphere.management.application.client.AppDeploymentController` instance to learn more about how to collect and populate tasks with WebSphere Application Server-specific binding information. The `AppDeploymentController` instance contains meta-data defined in XML-based deployment descriptors as well as annotations defined in Java classes within the input module.

```

//After you collect all the binding information, save it in the module.
controller.saveAndClose();

//Get the installation options.
Hashtable options = controller.getAppDeploymentSavedResults();

//Connect the administrative client, AdminClient, to WebSphere Application Server.
AdminClient client = ...;

//Create the application management proxy.
AppManagement proxy = AppManagementProxy.getJMXProxyForClient (client);

//Update the existing application, MyApp, by adding the module.
String appName = "MyApp";

options.put (AppConstants.APPUPDATE_CONTENTTYPE,
AppConstants. APPUPDATE_CONTENT_MODULEFILE);

//Create the notification filter.
NotificationFilterSupport myFilter = new NotificationFilterSupport();
myFilter.enableType (NotificationConstants.TYPE_APPMANAGEMENT);
//Add the listener.
NotificationListener listener = new AListener(_soapClient, myFilter,
"Install: " + appName, AppNotification.UPDATE);

//Specify the target for the new module.
Hashtable mod2svr = new Hashtable();
options.put (AppConstants.APPDEPL_MODULE_TO_SERVER, mod2svr);
mod2svr.put (uniquemoduleURI, target);
proxy.updateApplication ( appName,
moduleURI,
moduleName,
AppConstants.APPUPDATE_ADD,
options,
null);

// Wait for some timeout. The installation application programming interface (API) is
// asynchronous and so returns immediately.
// If the program does not wait here, the program ends.
Thread.sleep(300000); // Wait so that the program does not end.
}
catch (Exception e) {
e.printStackTrace();
}
}
}

// Specify the Java Management Extensions (JMX) notification listener for JMX events.
class AListener implements NotificationListener
{
AdminClient _soapClient;
NotificationFilterSupport myFilter;
Object handback;
ObjectName on;
String eventTypeToCheck;

public AListener(AdminClient cl, NotificationFilterSupport fl,
Object h, String eType) throws Exception
{
_soapClient = cl;
myFilter = fl;
handback = h;
eventTypeToCheck = eType;

Iterator iter = _soapClient.queryNames (new ObjectName(
"WebSphere:type=AppManagement,*"), null).iterator();
on = (ObjectName)iter.next();
System.out.println ("ObjectName: " + on);
_soapClient.addNotificationListener (on, this, myFilter, handback);
}

public void handleNotification (Notification notf, Object handback)
{
AppNotification ev = (AppNotification) notf.getUserData();
System.out.println ("!! JMX event Recd: (handback obj= " + handback+ "): " + ev);
}
}

```

```

//When the installation is done, remove the listener and quit
if (ev.taskName.equals (eventTypeToCheck) &&
    (ev.taskStatus.equals (AppNotification.STATUS_COMPLETED) ||
     ev.taskStatus.equals (AppNotification.STATUS_FAILED)))
{
    try
    {
        _soapClient.removeNotificationListener (on, this);
    }
    catch (Throwable th)
    {
        System.out.println ("Error removing listener: " + th);
    }
    System.exit (0);
}
}
}

```

What to do next

After you add the module successfully, remove the listener and quit.

Preparing and updating a module through programming

You can update a module for an existing application through the administrative console, the wsadmin tool, or programming. When you update a module, you replace the existing module with a new version. Use this example to update a module through programming.

Before you begin

This task assumes a basic familiarity with MBean programming. For information on MBean programming, see MBean Java application programming interface (API) documentation.

Before you can update a module on WebSphere Application Server, you must first install the application.

About this task

Perform the following tasks to update a module through programming.

Procedure

1. Create an application deployment controller instance to populate the Java archive file with binding information.
2. Save the binding information in the module.
3. Get the installation options.
4. If the preparation phase (population of the EAR file) is not performed, do the following actions:
 - a. Create an options table to be passed to the updateApplication MBean API.
 - b. Create a table for module to server relations and add the table to the options table.
5. Connect to WebSphere Application Server.
6. Create the application management proxy.
7. Create the notification filter.
8. Add the listener.
9. Replace the module in the application.
10. Specify the target for the new module.
11. Wait for some timeout so that the program does not end.
12. Listen to Java Management Extensions (JMX) notifications to understand completion of the operation.
13. When the module addition is done, remove the listener and quit.

Results

After you successfully run the code, the existing module is replaced with the new one.

Example

The following example shows how to add a module to an application based on the previous steps. Some statements are split on multiple lines for printing purposes.

```
//Inputs:
//moduleName specifies the name of the module that you add to the application.
//moduleURI specifies a URI that gives the target location of the module
// archive contents on a file system. The URI provides the location of the new
// module after installation. The URI is relative to the application URL.
//uniquemoduleURI specifies the URI that gives the target location of the
// deployment descriptor file. The URI is relative to the application URL.
//target specifies the cell, node, and server on which the module is installed.
//appName specifies the name of the application to update.

String moduleName = "/apps/foo.jar";
String moduleURI = "Increment.jar";
String uniquemoduleURI = "Increment.jar+META-INF/ejb-jar.xml";
String target = "WebSphere:cell=cellname,node=nodename,server=servername";
String appName = "MyApp";

//Get the administrative client to connect to
//WebSphere Application Server.
AdminClient client = ...;
AppManagement proxy = AppManagementProxy.getJMXProxyForClient (client);

Vector tasks = proxy.getApplicationInfo (appName, new Hashtable(), null);

//Create an application deployment controller instance, AppDeploymentController,
//to populate the Java archive (JAR) file with binding information.
//The binding information is WebSphere Application Server-specific deployment information.

Hashtable preferences = new Hashtable();
preferences.put (AppConstants.APPDEPL_LOCALE, Locale.getDefault());
preferences.put (AppConstants.APPUPDATE_CONTENTTYPE, AppConstants.APPUPDATE_CONTENT_MODULEFILE);
// When a module is update, there are three possible merge actions:
// AppConstants.APPREDEPL_DEFAULT_MERGE (the default merge action),
// This option specifies during the update action, the binding information from
// the new version of the EAR file or module is preferred
// over the corresponding binding information from the old version. If any element
// of binding is missing in the new version, the corresponding
// element from the old version is used.
// AppConstants.APPREDEPL_IGNORE_OLDBND,
// This option specifies that during the update action, the binding information
// from the new version of the module is preferred over
// the corresponding binding information from the old version. The bindings from
// the old version of the application or module are ignored.
// AppConstants.APPREDEPL_IGNORE_NEWBND
// This option specifies that during the update action, binding information from
// the old version of the module is preferred over
// the corresponding binding information from the new version. If any element of
// the binding does not exist in the old version,
// the element from the new version is used.
//
// To find the matching configuration object to perform the merge action, the values
// of the read only fields of the task for the new module are compared
// with the values of the read only fields of the task from the existing module.
// If all the read only values match, then the appropriate merge action is
// performed.
preferences.put (AppConstants.APPREDEPL_IGNORE_NEWBND, Boolean.TRUE);
AppDeploymentController controller = AppManagementFactory.readArchiveForUpdate(
    moduleName,
    moduleURI,
    AppConstants.APPUPDATE_UPDATE,
    preferences,
    tasks);
```

If the module that you update for the application lacks any bindings, add the bindings so that the module update works. Collect and add the bindings by using the public APIs that are provided with the product. Refer to Java documentation for the AppDeploymentController instance to learn more about how to collect and populate tasks with WebSphere Application Server-specific binding information. The AppDeploymentController instance contains meta-data defined in XML-based deployment descriptors as well as annotations defined in Java classes within the input module.

```
//After you collect all the binding information, save it in the module.
controller.saveAndClose();

//Create the notification filter.
```

```

NotificationFilterSupport myFilter = new NotificationFilterSupport();
myFilter.enableType (NotificationConstants.TYPE_APPMANAGEMENT);
//Add the listener.
NotificationListener listener = new AListener(_soapClient, myFilter,
"Install: " + appName, AppNotification.UPDATE);

//Get the installation options.
Hashtable options = controller.getAppDeploymentSavedResults();

//Update the existing application by adding the module.

options.put (AppConstants.APPUPDATE_CONTENTTYPE,
AppConstants. APPUPDATE_CONTENT_MODULEFILE);

//Specify the target for the new module
Hashtable mod2svr = new Hashtable();
options.put (AppConstants.APPDEPL_MODULE_TO_SERVER, mod2svr);
mod2svr.put (uniquemoduleURI, target);

proxy.updateApplication ( appName,
moduleURI,
moduleName,
AppConstants.APPUPDATE_UPDATE,
options,
null);
// Wait. The installation application programming interface (API) is
// asynchronous and so returns immediately.
// If the program does not wait here, the program ends.
Thread.sleep(300000); // Wait so that the program does not end.
}
catch (Exception e) {
e.printStackTrace();
}
}
}
// Specify the Java Management Extensions (JMX) notification listener for JMX events.
class AListener implements NotificationListener
{
AdminClient _soapClient;
NotificationFilterSupport myFilter;
Object handback;
ObjectName on;
String eventTypeToCheck;

public AListener(AdminClient cl, NotificationFilterSupport fl,
Object h, String eType) throws Exception
{
_soapClient = cl;
myFilter = fl;
handback = h;
eventTypeToCheck = eType;

Iterator iter = _soapClient.queryNames (new ObjectName(
"WebSphere:type=AppManagement,*"), null).iterator();
on = (ObjectName)iter.next();
System.out.println ("ObjectName: " + on);
_soapClient.addNotificationListener (on, this, myFilter, handback);
}

public void handleNotification (Notification notf, Object handback)
{
AppNotification ev = (AppNotification) notf.getUserData();
System.out.println ("!! JMX event Recd: (handback obj= " + handback+ "): " + ev);

//When the installation is done, remove the listener and quit

if (ev.taskName.equals (eventTypeToCheck) &&
(ev.taskStatus.equals (AppNotification.STATUS_COMPLETED) ||
ev.taskStatus.equals (AppNotification.STATUS_FAILED)))
{
try
{
_soapClient.removeNotificationListener (on, this);
}
catch (Throwable th)
{
System.out.println ("Error removing listener: " + th);
}
System.exit (0);
}
}
}
}

```

Adding a file through programming

You can add a file to an existing application through the administrative console, the wsadmin tool, or programming. This example describes how to add a file through programming.

Before you begin

This task assumes a basic familiarity with MBean programming. For information on MBean programming, see MBean Java application programming interface (API) documentation.

Before you can add a file to an application on WebSphere Application Server, you must install your application.

About this task

Perform the following tasks to add a file to an application through programming.

Procedure

1. Connect to WebSphere Application Server.
2. Create the application management proxy.
3. Create the notification filter for listening to events.
4. Add the listener.
5. Add the file to the application.
6. Wait for some timeout so that the program does not end.
7. Listen to Java Management Extensions (JMX) notifications to understand completion of the operation.
8. After you add a file, remove the listener and quit.

Results

After you successfully run the code, the file is added to the application.

Example

The following example shows how to add a file to an application based on the previous steps. Some statements are split on multiple lines for printing purposes.

```
import java.lang.*;
import java.io.*;
import java.util.*;
import java.lang.reflect.*;
import com.ibm.websphere.management.application.*;
import com.ibm.websphere.management.application.client.*;
import com.ibm.websphere.management.*;

import javax.management.*;

public class FileAdd {

    public static void main (String [] args) {

        try {

// Get a connection to WebSphere Application Server.
String host = "localhost";
String port = "8880";
String target = "WebSphere:cell=cellName,node=nodeName,server=server1";

Properties config = new Properties();
config.put (AdminClient.CONNECTOR_HOST, host);
config.put (AdminClient.CONNECTOR_PORT, port);
config.put (AdminClient.CONNECTOR_TYPE, AdminClient.CONNECTOR_TYPE_SOAP);
System.out.println ("Config: " + Config);
AdminClient _soapClient = AdminClientFactory.createAdminClient(config);

// Create the application management proxy, AppManagement.
AppManagement proxy = AppManagementProxy.getJMXProxyForClient (_soapClient);

String appName = "MyApp";
String fileURI = "test.war/com/acme/abc.jsp";

String fileContents = "/temp/abc.jsp";

//Create the notification filter.
NotificationFilterSupport myFilter = new NotificationFilterSupport();
myFilter.enableType (NotificationConstants.TYPE_APPMANAGEMENT);
```

```

//Add the listener.
NotificationListener listener = new AListener(_soapClient, myFilter,
"Install: " + appName, AppNotification.UPDATE);

Hashtable options = new Hashtable();
options.put (AppConstants.APPDEPL_LOCALE, Locale.getDefault());
options.put (AppConstants.APPUPDATE_CONTENTTYPE, AppConstants.APPUPDATE_CONTENT_FILE);

// Update the application
proxy.updateApplication ( appName,
    fileURI,
    fileContents,
    AppConstants.APPUPDATE_ADD,
    options,
    null);

// Wait; the installation Application Programming Interface (API) is
// asynchronous and so returns immediately.
// If the program does not wait here, the program ends.
Thread.sleep(90000); // Wait so that the program does not end.

    }
    catch (Exception e) {
        e.printStackTrace();
    }
}

// Specify the Java Management Extensions (JMX) notification listener for JMX events.
class AListener implements NotificationListener
{
    AdminClient _soapClient;
    NotificationFilterSupport myFilter;
    Object handback;
    ObjectName on;
    String eventTypeToCheck;

    public AListener(AdminClient cl, NotificationFilterSupport fl,
Object h, String eType) throws Exception
    {
        _soapClient = cl;
        myFilter = fl;
        handback = h;
        eventTypeToCheck = eType;

        Iterator iter = _soapClient.queryNames (new ObjectName(
"WebSphere:type=AppManagement,*"), null).iterator();
        on = (ObjectName)iter.next();
        System.out.println ("ObjectName: " + on);
        _soapClient.addNotificationListener (on, this, myFilter, handback);
    }

    public void handleNotification (Notification notf, Object handback)
    {
        AppNotification ev = (AppNotification) notf.getUserData();
        System.out.println ("!! JMX event Recd: (handback obj= " + handback+ "): " + ev);

        //When the installation is done, remove the listener and quit

        if (ev.taskName.equals (eventTypeToCheck) &&
            (ev.taskStatus.equals (AppNotification.STATUS_COMPLETED) ||
            ev.taskStatus.equals (AppNotification.STATUS_FAILED)))
        {
            try
            {
                _soapClient.removeNotificationListener (on, this);
            }
            catch (Throwable th)
            {
                System.out.println ("Error removing listener: " + th);
            }
            System.exit (0);
        }
    }
}

```

Updating a file through programming

You can update a file for an existing application through the administrative console, the wsadmin tool, or programming. This example describes how to update a file through programming.

Before you begin

This task assumes a basic familiarity with MBean programming. For information on MBean programming, see MBean Java application programming interface (API) documentation.

Before you can update a file for an application on WebSphere Application Server, you must install the application.

About this task

Perform the following tasks to update a file through programming.

Procedure

1. Connect to WebSphere Application Server.
2. Create the application management proxy.
3. Create the notification filter for listening to events.
4. Add the listener.
5. Update the file in the application.
6. Wait for some timeout so that the program does not end.
7. Listen to Java Management Extensions (JMX) notifications to understand completion of the operation.
8. When the installation completes, remove the listener and quit.

Results

After you successfully run the code, the file is updated for the application.

Example

The following example shows how to add a file to an application based on the previous steps. Some statements are split on multiple lines for printing purposes.

```
//Inputs:
//fileContents specifies the name of the file that you add to the application.
//appName specifies the name of the application.
//fileURI specifies a URI that gives the target location of the file. The URI
// provides the location of the new module after installation. The URI is
// relative to the application URL.

String fileContents = "/apps/test.jsp";
String appName = "MyApp";
String fileURI = "SomeWebMod.war/com/foo/abc.jsp";

//Get the administrative client to connect to
//WebSphere Application Server.
AdminClient client = ...;

//Create the application management proxy.
AppManagement proxy = AppManagementProxy.getJMXProxyForClient (client);

//Create the notification filter.
NotificationFilterSupport myFilter = new NotificationFilterSupport();
myFilter.enableType (NotificationConstants.TYPE_APPMANAGEMENT);
//Add the listener.
NotificationListener listener = new AListener(_soapClient, myFilter,
"Install: " + appName, AppNotification.UPDATE);

Hashtable options = new Hashtable();
options.put (AppConstants.APPDEPL_LOCALE, Locale.getDefault());
options.put (AppConstants.APPUPDATE_CONTENTTYPE, AppConstants.APPUPDATE_CONTENT_FILE);

proxy.updateApplication ( appName,
    fileURI,
    fileContents,
    AppConstants.APPUPDATE_UPDATE,
    options,
    null);

// Wait; the installation application programming interface (API) is
// asynchronous and so returns immediately.
```

```

// If the program does not wait here, the program ends.
Thread.sleep(300000); // Wait so that the program does not end.
    }
    catch (Exception e) {
        e.printStackTrace();
    }
}
}
// Specify the Java Management Extensions (JMX) notification listener for JMX events.
class AListener implements NotificationListener
{
    AdminClient _soapClient;
    NotificationFilterSupport myFilter;
    Object handback;
    ObjectName on;
    String eventTypeToCheck;

    public AListener(AdminClient cl, NotificationFilterSupport fl,
Object h, String eType) throws Exception
    {
        _soapClient = cl;
        myFilter = fl;
        handback = h;
        eventTypeToCheck = eType;

        Iterator iter = _soapClient.queryNames (new ObjectName(
"WebSphere:type=AppManagement,*"), null).iterator();
        on = (ObjectName)iter.next();
        System.out.println ("ObjectName: " + on);
        _soapClient.addNotificationListener (on, this, myFilter, handback);
    }

    public void handleNotification (Notification notf, Object handback)
    {
        AppNotification ev = (AppNotification) notf.getUserData();
        System.out.println ("!! JMX event Recd: (handback obj= " + handback+ "): " + ev);

        //When the installation is done, remove the listener and quit.

        if (ev.taskName.equals (eventTypeToCheck) &&
            (ev.taskStatus.equals (AppNotification.STATUS_COMPLETED) ||
             ev.taskStatus.equals (AppNotification.STATUS_FAILED)))
        {
            try
            {
                _soapClient.removeNotificationListener (on, this);
            }
            catch (Throwable th)
            {
                System.out.println ("Error removing listener: " + th);
            }
            System.exit (0);
        }
    }
}
}

```

Uninstalling an application through programming

You can uninstall an application through the administrative console, the wsadmin tool, or programming. Use this example to uninstall an application through programming.

Before you begin

This task assumes a basic familiarity with MBean programming. For information on MBean programming, see MBean Java application programming interface (API) documentation.

Before you can uninstall an application on WebSphere Application Server, you must first install it.

About this task

Perform the following tasks to uninstall an application through programming.

Procedure

1. Get a connection to WebSphere Application Server.
2. Get the application management proxy.
3. Create the notification filter for listening to uninstallation events.

4. Add the listener.
5. Uninstall the application.
6. Wait for some timeout so that the program does not end.
7. Listen to Java Management Extensions (JMX) notifications to understand completion of the operation.
8. When the uninstallation completes, remove the listener and quit.

Results

After you successfully run the code, the application is uninstalled.

Example

The following example shows how to uninstall an application based on the previous steps. Some statements are split on multiple lines for printing purposes.

```
import java.lang.*;
import java.io.*;
import java.util.*;
import java.lang.reflect.*;
import com.ibm.websphere.management.application.*;
import com.ibm.websphere.management.application.client.*;
import com.ibm.websphere.management.*;

import javax.management.*;

public class Uninstall {

    public static void main (String [] args) {

        try {

// Get a connection to the server.
String host = "localhost";
String port = "8880";
String target = "WebSphere:cell=cellName,node=nodeName,server=server1";

Properties config = new Properties();
config.put (AdminClient.CONNECTOR_HOST, host);
config.put (AdminClient.CONNECTOR_PORT, port);
config.put (AdminClient.CONNECTOR_TYPE, AdminClient.CONNECTOR_TYPE_SOAP);
System.out.println ("Config: " + config);
AdminClient _soapClient = AdminClientFactory.createAdminClient(config);

// Get the application management proxy.
AppManagement proxy = AppManagementProxy.getJMXProxyForClient (_soapClient);

String appName = "MyApp";
Hashtable options = new Hashtable();
options.put (AppConstants.APPDEPL_LOCALE, Locale.getDefault());

//Create the notification filter.
NotificationFilterSupport myFilter = new NotificationFilterSupport();
myFilter.enableType (AppConstants.NotificationType);

//Add the listener.
NotificationListener listener = new AListener(_soapClient,
myFilter, "Install: " + appName, AppNotification.UNINSTALL);

// Uninstall the application.
proxy.uninstallApplication (appName, options, null);
System.out.println ("After uninstall App is called..");

// Wait for some timeout. The installation application programming interface (API) is
// asynchronous and so returns immediately.
// If the program does not wait here, the program ends.
Thread.sleep(300000); // Wait so that the program does not end.

        }
        catch (Exception e) {
            e.printStackTrace();
        }
    }
}

// Specify the Java Management Extensions (JMX) notification listener for JMX events.
class AListener implements NotificationListener
{
    AdminClient _soapClient;
    NotificationFilterSupport myFilter;
    Object handback;
}
```

```

ObjectName on;
String eventTypeToCheck;

public AListener(AdminClient cl, NotificationFilterSupport fl,
Object h, String eType) throws Exception
{
    _soapClient = cl;
    myFilter = fl;
    handback = h;
    eventTypeToCheck = eType;

    Iterator iter = _soapClient.queryNames (new ObjectName(
"WebSphere:type=AppManagement,*"), null).iterator();
    on = (ObjectName)iter.next();
    System.out.println ("ObjectName: " + on);
    _soapClient.addNotificationListener (on, this, myFilter, handback);
}

public void handleNotification (Notification notf, Object handback)
{
    AppNotification ev = (AppNotification) notf.getUserData();
    System.out.println ("!! JMX event Recd: (handback obj= " + handback+ "): " + ev);

    //When the unistallation is done, remove the listener and quit

    if (ev.taskName.equals (eventTypeToCheck) &&
        (ev.taskStatus.equals (AppNotification.STATUS_COMPLETED) ||
         ev.taskStatus.equals (AppNotification.STATUS_FAILED)))
    {
        try
        {
            _soapClient.removeNotificationListener (on, this);
        }
        catch (Throwable th)
        {
            System.out.println ("Error removing listener: " + th);
        }
        System.exit (0);
    }
}
}
}

```

Deleting a module through programming

You can delete a module from an existing application through the administrative console, the wsadmin tool, or programming. Use this example to delete a module through programming.

Before you begin

This task assumes a basic familiarity with MBean programming. For information on MBean programming, see MBean Java application programming interface (API) documentation.

Before you can delete a module from an application on WebSphere Application Server, you must install the application.

About this task

Perform the following tasks to delete a module through programming.

Procedure

1. Connect to WebSphere Application Server.
2. Create the application management proxy.
3. Create the notification filter for listening to events.
4. Add the listener.
5. Delete the module.
6. Wait for some timeout so that the program does not end.
7. Listen to Java Management Extensions (JMX) notifications to understand completion of the operation.
8. When the module is deleted, remove the listener and quit.

Results

After you successfully run the code, the existing module is deleted from the application.

Example

The following example shows how to delete a module from an application based on the previous steps. Some statements are split on multiple lines for printing purposes.

```
//moduleURI specifies a URI that gives the target location of the module.
//appName specifies the name of the application to update.
String moduleURI = "Increment.jar";
String appName = "MyApp";

//Get the administrative client to connect to
//WebSphere Application Server.
AdminClient client = ...;

//Create the application management proxy.
AppManagement proxy = AppManagementProxy.getJMXProxyForClient (client);

//Create the notification filter.
NotificationFilterSupport myFilter = new NotificationFilterSupport();
myFilter.enableType (NotificationConstants.TYPE_APPMANAGEMENT);
//Add the listener.
NotificationListener listener = new AListener(_soapClient, myFilter,
"Install: " + appName, AppNotification.UPDATE);

//Update the existing application, MyApp, by deleting the module.
Hashtable options = new Hashtable();
options.put (AppConstants.APPDEPL_LOCALE, Locale.getDefault());
options.put (AppConstants.APPUPDATE_CONTENTTYPE, AppConstants.APPUPDATE_CONTENT_MODULEFILE);

proxy.updateApplication ( appName,
    moduleURI,
    null,
    AppConstants.APPUPDATE_DELETE,
    options,
    null);

// Wait; the installation application programming interface (API) is
// asynchronous and so returns immediately.
// If the program does not wait here, the program ends.
Thread.sleep(300000); // Wait so that the program does not end.
    }
    catch (Exception e) {
        e.printStackTrace();
    }
}
}

// Specify the Java Management Extensions (JMX) notification listener for JMX events.
class AListener implements NotificationListener
{
    AdminClient _soapClient;
    NotificationFilterSupport myFilter;
    Object handback;
    ObjectName on;
    String eventTypeToCheck;

    public AListener(AdminClient cl, NotificationFilterSupport fl,
Object h, String eType) throws Exception
    {
        _soapClient = cl;
        myFilter = fl;
        handback = h;
        eventTypeToCheck = eType;

        Iterator iter = _soapClient.queryNames (new ObjectName(
"WebSphere:type=AppManagement,*"), null).iterator();
        on = (ObjectName)iter.next();
        System.out.println ("ObjectName: " + on);
        _soapClient.addNotificationListener (on, this, myFilter, handback);
    }

    public void handleNotification (Notification notf, Object handback)
    {
        AppNotification ev = (AppNotification) notf.getUserData();
        System.out.println ("!! JMX event Recd: (handback obj= " + handback+ "): " + ev);

        //When the installation is done, remove the listener and quit

        if (ev.taskName.equals (eventTypeToCheck) &&
(ev.taskStatus.equals (AppNotification.STATUS_COMPLETED) ||
ev.taskStatus.equals (AppNotification.STATUS_FAILED)))
        {

```

```

        try
        {
            _soapClient.removeNotificationListener (on, this);
        }
    catch (Throwable th)
    {
        System.out.println ("Error removing listener: " + th);
    }
    System.exit (0);
}
}
}

```

Deleting a file through programming

You can delete a file from an existing application through the administrative console, the wsadmin tool, or programming. Use this example to delete a file through programming.

Before you begin

This task assumes a basic familiarity with MBean programming. For information on MBean programming, see MBean Java application programming interface (API) documentation.

Before you can delete a file from an application on WebSphere Application Server, you must install the application.

About this task

Perform the following tasks to delete a file through programming.

Procedure

1. Connect to WebSphere Application Server.
2. Create the application management proxy.
3. Create the notification filter for listening to events.
4. Add the listener.
5. Delete the file from the application.
6. Wait for some timeout so that the program does not end.
7. Listen to Java Management Extensions (JMX) notifications to understand completion of the operation.
8. When the file is deleted from the application, remove the listener and quit.

Results

After you successfully run the code, the file is deleted from the application.

Example

The following example shows how to delete a file based on the previous steps. Some statements are split on multiple lines for printing purposes.

```

//Inputs:
//fileURI specifies a URI that gives the target location of the file. The URI
// provides the location of the new module after installation. The URI is
// relative to the application URL.
//appName specifies the name of the application.

String fileURI = "Increment.jar/com/acme/Foo.class";
String appName = "MyApp";

//Get the administrative client to connect to
//WebSphere Application Server.
AdminClient client = ...;

//Create the application management proxy.
AppManagement proxy = AppManagementProxy.getJMXProxyForClient (client);

```



```

//Create the notification filter.
NotificationFilterSupport myFilter = new NotificationFilterSupport();
myFilter.enableType (NotificationConstants.TYPE_APPMANAGEMENT);
//Add the listener.
NotificationListener listener = new AListener(_soapClient, myFilter,
"Install: " + appName, AppNotification.UPDATE);

//Update the existing application, MyApp, by deleting the file.
Hashtable options = new Hashtable();
options.put (AppConstants.APPDEPL_LOCALE, Locale.getDefault());
options.put (AppConstants.APPUPDATE_CONTENTTYPE, AppConstants.APPUPDATE_CONTENT_FILE);

proxy.updateApplication ( appName,
fileURI,
null,
AppConstants.APPUPDATE_DELETE,
options,
null);

// Wait for some timeout. The installation Application Programming Interface (API) is
// asynchronous and so returns immediately.
// If the program does not wait here, the program ends.
Thread.sleep(300000); // Wait so that the program does not end.
    }
    catch (Exception e) {
        e.printStackTrace();
    }
}

}

// Specify the Java Management Extensions (JMX) notification listener for JMX events.
class AListener implements NotificationListener
{
    AdminClient _soapClient;
    NotificationFilterSupport myFilter;
    Object handback;
    ObjectName on;
    String eventTypeToCheck;

    public AListener(AdminClient cl, NotificationFilterSupport fl,
Object h, String eType) throws Exception
    {
        _soapClient = cl;
        myFilter = fl;
        handback = h;
        eventTypeToCheck = eType;

        Iterator iter = _soapClient.queryNames (new ObjectName(
"WebSphere:type=AppManagement,*"), null).iterator();
        on = (ObjectName)iter.next();
        System.out.println ("ObjectName: " + on);
        _soapClient.addNotificationListener (on, this, myFilter, handback);
    }

    public void handleNotification (Notification notf, Object handback)
    {
        AppNotification ev = (AppNotification) notf.getUserData();
        System.out.println ("!! JMX event Recd: (handback obj= " + handback+ "): " + ev);

        //Once the installation is done, remove the listener and quit

        if (ev.taskName.equals (eventTypeToCheck) &&
(ev.taskStatus.equals (AppNotification.STATUS_COMPLETED) ||
ev.taskStatus.equals (AppNotification.STATUS_FAILED)))
        {
            try
            {
                _soapClient.removeNotificationListener (on, this);
            }
            catch (Throwable th)
            {
                System.out.println ("Error removing listener: " + th);
            }
            System.exit (0);
        }
    }
}
}

```

Chapter 10. Extending application management operations through programming

You can use the common deployment framework to add additional logic to application management operations. The additional logic can do such tasks as code generation, configuration operations, additional validation, and so on. This topic demonstrates, through programming, how to plug into the *common deployment framework* to extend application management operations.

Before you begin

This task assumes a basic familiarity with Java application programming interfaces (APIs). Read about the Java APIs in the application programming interfaces documentation.

Before you can extend application management operations, you must first install WebSphere Application Server.

About this task

Use this example to extend application management through programming. The tasks that the extensions provide are available through all the administrative clients, such as the wsadmin tool, the administrative console, or through programmatic APIs that the AppManagement MBean provides.

Procedure

1. Define your extension as an Eclipse plug-in and add a `plugin.xml` file to register your extension provider with the deployment framework.
 - a. In the `plugin.xml` file, provide an extension provider implementation class for the `common-deployment-framework-extensionprovider` extension point.
 - b. Put the plug-in Java archive (JAR) file in the `plugins` directory of your WebSphere Application Server installation.

```
<?xml version="1.0" encoding="UTF-8"?>
<plugin
  id="com.ibm.myproduct.MyExtensionProvider"
  name="My Extension"
  version="1.0.0">

  <extension point="common-deployment-framework-extensionprovider">
    <action class="com.acme.MyExtendProviderImpl"/>
  </extension>
</plugin>
```

2. Provide an extension provider.

An extension provider class provides steps for a given operation on an application Enterprise archive (EAR) file. Before an operation runs, the deployment framework queries all the registered extension providers for additional steps. A single list of steps is passed to each provider. Each provider can add steps to the list. The default provider that the deployment framework provides is called first to populate the list with default steps. Other extension providers are called next.

Various operations that you can extend through the common deployment framework are defined as constants in the `DeploymentConstants` class. These operations are described in the following table. Some operations are split on multiple lines for printing purposes.

Table 41. Extensible `DeploymentConstants` operations. Select an operation to extend..

Operation	Description
<code>DeploymentConstants.CDF_OP_INSTALLJ2EE</code>	Installs a Java Platform, Enterprise Edition (Java EE) EAR file
<code>DeploymentConstants.CDF_OP_EDITJ2EE</code>	Edits a deployment application configuration

Table 41. Extensible DeploymentConstants operations (continued). Select an operation to extend..

Operation	Description
DeploymentConstants.CDF_OP_UPDATEJ2EE	Applies a fine-grained update to an application such as addition, removal, or update of a file or a module; or partial update of an application
DeploymentConstants.CDF_OP_UNINSTALLJ2EE	Uninstalls a Java EE application
DeploymentConstants. CDF_OP_CREATE_EAR_WRAPPERJ2EE	Wraps the contents input to the application installation into an EAR file

The AppManagement MBean, which is responsible for deploying and managing Java EE applications on WebSphere Application Server, runs all the operations except the CDF_OP_CREATE_EAR_WRAPPERJ2EE operation. Deploy the extensions that extend these operations in the plugins directory of the deployment manager.

Either the wsadmin utility or the administrative console runs the CDF_OP_CREATE_EAR_WRAPPERJ2EE operation when the input contents that are supplied to the CDF_OP_INSTALLJ2EE operation are not packaged as an EAR file. Deploy an extension that extends the CDF_OP_CREATE_EAR_WRAPPERJ2EE operation in the plugins directory of the wsadmin installation.

The following example provides an extension provider that does the following tasks:

- a. Adds two additional steps for the application installation operation
- b. Adds one step for wrapping input contents into an EAR file

```
package com.acme;

import com.ibm.websphere.management.deployment.registry.ExtensionProvider;
import com.ibm.websphere.management.deployment.core.DeploymentConstants;

public class MyExtensionProviderImpl extends ExtensionProvider {
    public void addSteps (String type, String op, String phase,
        List steps)
    {
        if (op.equals (DeploymentConstants.CDF_OP_INSTALLJ2EE))
        {
            // Add a code generation step.
            steps.add (0, new com.acme.CodeGenStep());
            // Add a configuration step.
            steps.add (new com.acme.ConfigStep());
        }
        else if (op.equals (DeploymentConstants.CDF_OP_CREATE_EAR_WRAPPERJ2EE))
        {
            // Add an ear-wrapper step.
            steps.add (new com.acme.EarWrapperStep());
        }
    }
}
```

3. Provide the deployment step implementation.

An extension provider adds a deployment step. The step contains logic that performs additional processing in an application management operation. The logic provides the step access to the deployment context and the deployable object. The deployment context provides information, such as the name of the operation, the configuration session ID, a temporary location for creating temporary files, operation parameters, and so on. The deployable object wraps the deployment content input to the operation. For example, the deployable object wraps the Java EE EAR file for the installation operation or a file, a module, or a partial application for the update operation.

- The following example illustrates how an extension during installation entirely changes an EAR file that is input to the installation operation. The example provides a deployment step during the installation operation that does the following tasks:
 - a. Runs code generation to generate a new EAR file.
 - b. Calls the setContentPath method in the DeployableObject class to set the new EAR file path. The default installation logic, such as steps that the default installation logic adds, uses this new EAR file for installation in the configuration repository.

```

package com.acme;

import com.ibm.websphere.management.deployment.core.DeploymentStep;
import com.ibm.websphere.management.deployment.core.DeployableObject;

public class CodeGenStep extends DeploymentStep
{
    public void execute (DeployableObject dObject)
    {
        EARFile earFile = (EARFile)dObject.getHandle();
        String newEARPath = null;
        // Use step specific logic to create another EAR file after code generation.
        ...
        newEARPath = _context.getTempDir() + "new.ear";

        dObject.setContentPath (newEARPath);
    }
}

```

- The following example provides a deployment step that:
 - a. Reads the contents of the input EAR file.
 - b. Manipulates the configuration session accessed through the context instance, `_context`.

```

package com.acme;

public class ConfigStep extends DeploymentStep
{
    public void execute (DeployableObject dObject)
    {
        EARFile earFile = (EARFile) dObject.getHandle();

        // Use the following example code to perform the configuration.
        String sessionID = _context.getSessionID();
        com.ibm.websphere.management.Session session = new
            com.ibm.websphere.management.Session (sessionID, true);
        // Use the configuration service to perform the configuration steps.
        ...

        // Read the application configuration.
        Application appDD = earFile.getDeploymentDescriptor();
        ...

        String newEARPath = null;
    }
}

```

- The following example provides a deployment step to wrap arbitrary content around an EAR file. Application management logic accepts only the EAR file for deployment. An extension is required if you want to input anything other than an EAR file to the deployment process.

```

package com.acme;

import com.ibm.websphere.management.deployment.core.DeploymentStep;
import com.ibm.websphere.management.deployment.core.DeployableObject;

public class EarWrapperStep extends DeploymentStep
{
    public void execute (DeployableObject dObject)
    {
        Archive archive = (Archive) dObject.getHandle();
        String newEARPath = null;
        // provide your logic to wrap the jar with the ear
        ...
        newEARPath = //;
        // Set the new ear path back into DeploymentContext
        this.getContext().getContextData()
            .put(DeploymentContext.RETURN_Object_key, newEARPath);
    }
}

```

Results

Through programming, you have plugged into the common deployment framework to extend application management operations.

What to do next

You can extend other application management operations, or do any other administrative operations you choose.

Chapter 11. Deploying and administering business-level applications

Deploying a business-level application consists of creating the business-level application on a Version 7.0 or later server.

Before you begin

A business-level application is an administration model that provides the entire definition of an application as it makes sense to the business. It is a WebSphere configuration artifact, similar to a server or cluster, that is stored in the product configuration repository. A business-level application can contain artifacts such as Java Platform, Enterprise Edition (Java EE) applications or modules, shared libraries, data files, and other business-level applications. You might use a business-level application to group related artifacts or to add capability to an existing application. For example, suppose you want to add capability provided in a Java archive (JAR) to a Java EE application already deployed on a product server. You can add that capability by creating a new business-level application and adding the JAR file and the deployed Java EE application to the business-level application. In some cases, you do not even need to change the deployed Java EE application configuration to add the capability.

Before creating a business-level application, you must develop the artifacts to go in the application and configure the target server or cluster. Before choosing a deployment target for the application, ensure that the target version is 7.0 or later.

About this task

When creating a business-level application, you can configure the application enough to enable it to run on the server. Later, you can configure the application and its contents further, start or stop the application, and otherwise manage its activity.

The topics in this section describe how to deploy and administer a business-level application or its contents using the administrative console. You can also use programming or wsadmin scripting.

Procedure

- Import assets to a repository.
- View, delete, update, or export assets.
- Create a business-level application.
- Create a Service Component Architecture (SCA) business-level application that has SCA assets, shared libraries, or business-level applications.
- Start the application.
- Stop the application.
- Update the application and its configuration units.
- Update SCA composite artifacts.
- View the composite definition of an SCA asset composition unit.
- View SCA domain information.
- “Viewing and editing JMS bindings on references and services of SCA composites” on page 463
- Delete the application.

What to do next

After making changes to administrative configurations of your applications in the administrative console, ensure that you save the changes.

Business-level applications

A business-level application is an administration model that provides the entire definition of an application as it makes sense to the business. A business-level application is a WebSphere configuration artifact, similar to a server or cluster, that is stored in the product configuration repository.


- Business-level application characteristics
- Comparisons to Java EE applications

Business-level application characteristics

A business-level application has the following characteristics:

- A business-level application is an administration model of the definition of an enterprise-level application that consists of WebSphere and non-WebSphere artifacts. The business-level application might not explicitly manage the lifecycle of every artifact. It is a model for defining an application.
- A business-level application does not represent or contain application binary files. It is a configuration that lists one or more composition units, which represent the application binary files. A business-level application uses the binary files to run the application business logic. Administration of binary files is separate from administration of the application definition.
- A business-level application supports recursive composition by reference that facilitates hierarchical assembly of business-level applications and individual deployed artifacts within or outside a WebSphere product. The composition at its lowest level consists of configured instances of application binary files that run in a specific runtime environment such as an application server. Installable packages or archives, such as Java archives (JAR) or enterprise archive (EAR) files, typically deliver the business logic that these configured instances represent to corresponding runtime platforms.

The following diagram shows the composition model for business-level applications:



A business-level application does not introduce new programming, runtime, or packaging models:

- You do not need to change your application business logic. The business-level application function does not introduce new application programming interfaces (APIs).
- You do not need to change your application runtime settings. The product supports all of the runtime characteristics, such as security, class loading and isolation, required by individual programming models to which business components are written.
- You do not need to change your application packaging. There is no specific unique packaging model that provides a business-level application definition.

Typically, you first create an empty business-level application and then add composition units to it. The business-level application name must be unique within a cell. The business level application itself has minimal configuration data associated with it, solely the list of composition units, but individual composition units might save application-specific configuration data.

A business-level application is defined in the product configuration repository under *profile_root/config/cells/cell_name/b1as/business_level_application_name/bver/BASE/b1a.xml*.

Comparisons to Java EE applications

Business-level applications can consist of or aggregate Java Platform, Enterprise Edition (Java EE) applications and modules with non-Java EE artifacts. The contents of Java EE applications integrate with business-level application concepts for deployment and management of applications. Existing Java EE application management APIs continue to work after you add Java EE application or modules to a business-level application. The business-level application management API accepts Java EE contents and configurations and delegates to existing Java EE management APIs. Control operations such as starting and stopping a Java EE composition unit are delegated to ApplicationManager MBean on application servers that start and stop Java EE applications.

Table 42. Java EE concepts compared to business-level application concepts. Business-level application concepts include assets, composition units, and deployable units.

Java EE concept	Business-level application concept	Description
EAR or stand-alone module for deployment	Asset	Java EE application contents are assets.
Java EE application created at the end of application install	Composition unit	A Java EE application is in an enterprise archive (EAR) file. The product saves the EAR file in the product repository as a composition unit.
Java EE modules within the EAR file	Deployable units in the asset	Each module in the EAR file is a deployable unit that you can install on independent deployment targets. The EAR file is still managed as a single asset in its entirety.

Table 42. Java EE concepts compared to business-level application concepts (continued). Business-level application concepts include assets, composition units, and deployable units.

Java EE concept	Business-level application concept	Description
Java EE application installation using the administrative console, programming, or wsadmin commands	<p>Multiple business-level application management commands</p> <p>During Java EE application deployment, you can specify the name of the business-level application to include the Java EE application. If the business-level application name is not set, the product creates a default business-level application with the same name as the Java EE application name. The product adds a composition unit with the same name as the Java EE application name under the business-level application. You can deploy multiple Java EE applications under a single business-level application.</p>	<p>You can make a Java EE application a business-level application and add it to another business-level application:</p> <ol style="list-style-type: none"> 1. Install the Java EE application (EAR file) using the enterprise application installation console wizard, programming, or wsadmin. Keep the default selection to create a business-level application that has the same name as the Java EE application. 2. Create an empty business-level application. 3. Add the EAR file business-level application to the empty business-level application. The EAR file business-level application is a composition unit of the containing business-level application. <p>Or, you can make a Java EE application an asset and add it to another business-level application:</p> <ol style="list-style-type: none"> 1. Import an EAR file as an asset. It has an asset type aspect of Java EE ear. 2. Create an empty business-level application. 3. Add the Java EE application asset to the business-level application. The EAR file asset is a composition unit of the containing business-level application. 4. Collect targets for each deployable unit (Java EE module).
Uninstall Java EE application	Multiple business-level application management commands	<p>You delete the Java EE application composition unit from the business-level application:</p> <ol style="list-style-type: none"> 1. Remove the composition unit for the Java EE application from the business-level application. 2. If the EAR file is an asset, delete the asset.
Start the Java EE application.	Start the composition unit.	Starting a business-level application starts any Java EE application in it.
Stop the Java EE application.	Stop the composition unit.	Stopping a business-level application stops any Java EE application in it.

Assets

An asset represents one or more application binary files that are stored in an asset repository. Typical assets include application business logic such as Java Platform, Enterprise Edition (Java EE) archives, library files, and other resource files.

An asset repository stores the binary files for the asset. The product configuration repository provides a default asset repository.

Assets in the configuration repository are managed by the product management domain. The configuration repository stores asset binary files in `app_server_root/config/cells/cell_name/assets/asset_name/aver/BASE/bin/`.

An asset name must be unique within a cell, the product administrative domain.

The product creates an `asset.xml` file when an asset is registered with the product configuration. The file contains information about the asset such as its name, destination location, and dependencies on other assets.

You must register files as assets before you can add them to one or more business-level applications. At the time of asset registration, you can import the physical application files into the product configuration repository or you can specify an external location where the asset resides.

Composition units

A composition unit represents a configured asset in a business-level application. A composition unit enables the asset contents to interact with other assets in the application. It also enables the product run time to load and run asset contents.

The product supports three types of composition units:

Asset composition units

Composition units created from assets by configuring each deployable unit of the asset to run on deployment targets.

Shared library composition units

Composition units created from JAR-based assets by ignoring all the deployable objects from the asset and treating the asset JAR file as a library of classes.

Business-level application composition units

Composition units created from business-level applications that are added to existing business-level applications.

A composition unit contains the following information:

- Configuration information that binds contents of an asset with a specific hosting run time and adds the configuration necessary for the run time to load and run the asset
- References to external services, components, or other resources that the asset uses
- Customized configurations for service definitions, references and other relevant configuration data
- A list of deployment targets or runtime environments along with the runtime environment-specific configuration where the composition unit runs.

For example, a composition unit for an enterprise bean (EJB) Java archive (JAR) asset is an EJB module instance that contains necessary EJB binding information, such as EJB Java Naming and Directory Interface (JNDI) names and `ejb-ref` resolutions, along with a list of application servers or clusters where the EJB JAR runs.


The product creates a composition unit from only one asset. However, multiple composition units can share a single asset. This is particularly useful in scenarios where different configurations use the same application binary files to provide different runtime behavior.

The following rules apply to a composition unit:

- A composition unit can exist only in a business-level application.
- Because a composition unit contains application-specific configuration and wiring information, multiple business-level applications cannot share an asset or shared library composition unit.

The following graphic shows the use of composition units in business-level applications. Assume that you have unprocessed files, such as archives, that you want to use in business-level applications. Before you can add the files to business-level applications, you must first import the files as assets, which adds the files to the product repository. Next, you add the assets to business-level applications, which creates composition units for the assets. Business-level applications can contain asset composition units, shared

library composition units, or business-level composition units.



Importing assets

You must register application business logic such as Java Platform, Enterprise Edition (Java EE) archives, libraries, and other resource files with the product configuration as assets before you can add the assets to one or more business-level applications. Importing an asset registers it with the product configuration.

Before you begin

This topic assumes that you have one or more application binary files that you want to add to a business-level application. You must register those binary files as assets before you can add them to the business-level application.

About this task

Before a business-level application that uses an asset can be started on the target run time, the asset binaries must be extracted to a deployer-defined location on the file system that is local to the target run time. Importing an asset extracts binaries to a location that is local to the target run time.

The application server run time that reads the asset binaries either at application start time or while serving an incoming client request determines the extraction format of the asset binaries. The extraction format might include unzipping of Java archive (JAR) or compressed (zip) files.

This topic describes how to import an asset using the administrative console. Alternatively, you can use the wsadmin tool or programming.

Procedure

1. Click **Applications > New Application > New Asset** in the console navigation tree.
2. On the Upload asset page, specify the asset package to import.
 - a. Specify the full path name of the asset.
 - b. Click **Next**.
3. On the Select options for importing an asset page, specify asset settings.

You typically can click **Next** and use the default values.

 - a. Optional: For **Asset description**, specify a brief description of the asset.
 - b. Optional: For **Asset binaries destination URL**, specify the target location of the asset.

This setting specifies the location to which the product extracts the asset. After an asset is imported, the product looks for the asset in this location when a running application uses the asset.

If you do not specify a value, the product installs the asset to the default location, `${profile_root}/installedAssets/asset_name/BASE/`.
 - c. Optional: For **Asset type aspects**, examine the asset content type and version specified by the product. You cannot change this setting value.

The type aspect typically denotes the type of application contents, such as a specification to which the application is written. For example, an enterprise bean (EJB) that supports the EJB Version 2.0 specification has the aspects `type=EJB,version=2.0`.

If the type aspect is none and if the asset is a JAR file, then the product associates a `javarchive` type aspect with the asset by default.
 - d. For **File permissions**, specify any file permissions that are set on asset binary files so the target run time can read or run the asset. Importing the asset extracts its binary files on the disk local to the target runtime environment.

Try importing the asset using the default value. For detailed information on the **File permissions** setting, refer to the Select options for importing an asset page online help.
 - e. For **Current asset relationships**, add assets that the asset you are importing needs to run or remove assets that are not needed.

When the product imports a JAR asset, the product detects asset relationships automatically by matching the dependencies defined in the JAR manifest with the assets that are already imported into the administrative domain.
 - f. For **Validate asset**, specify whether the product validates the asset.

The setting is deselected by default. This **false (no)** value is appropriate for most assets. Only select **true (yes)** to validate an asset when needed.

The product does not save the value specified for **Validate asset**. Thus, if you select to validate the asset (**yes**) now and later update the asset, when you update the asset you must enable this setting again for the product to validate the updated files.
 - g. Click **Next**.
4. On the Summary page, click **Finish**.

Results

Several messages are displayed, indicating whether your asset is imported successfully.

An asset can contain multiple deployable objects as defined by the application contents of that asset. A *deployable object* is a part of the asset that you can map to a deployment target such as an application server or a cluster. If the product imports the asset successfully, then appropriate deployable objects are identified in the asset and are further used when a composition unit is created from that asset.

If the asset importing is not successful, read the messages and try importing the asset again. Correct the values noted in the messages.

What to do next

If the product imports the asset successfully and displays the list of assets on the Assets page, then click **Save**.

Add a composition unit to a business-level application using the asset that you imported. An asset included in a business-level application is represented by a composition unit.

Upload asset settings

Use this page to specify the asset to register with the asset repository. You can add registered assets to a business-level application.

To view this administrative console page, click **Applications > New application > New Asset**.

Importing an asset registers the asset with the asset repository.

The product manages the contents of a registered asset as a single entity. The contents of a registered asset must be accessible to application servers, web servers and other runtime environments that use the asset.

During asset importing, asset files typically are uploaded from a client workstation running the browser to the server running the administrative console, where they are registered. In such cases, use the web browser running the administrative console to select files to upload to the server.

Path to the asset

Specifies the fully qualified path to the asset.

Specify one of the following supported assets:

- A single file, such as an enterprise bean (EJB) file
- An archive of files, such as a Java archive (JAR) or a compressed .zip file
- An archive of archives, such as an enterprise archive (EAR) or shared library file

Use **Local file system** if the browser and asset files are on the same machine (whether or not the server is on that machine, too).

Use **Remote file system** if the asset file resides on any node in the current cell context. Only supported assets are shown during the browsing. Also use **Remote file system** to specify an asset file that is already residing on the machine running the application server. For example, the field value might be *profile_root/installableApps/my_bean.ejb*. After the asset file is transferred, the **Remote file system** value shows the path of the temporary location on the server.

Asset settings

Use this page to specify options for the registration of an asset with the asset repository. Default values for the options are used if you do not specify a value. If the asset is an OSGi application, additional information about bundle download status is displayed.

To view this administrative console page, click **Applications > Application Types > Assets > *asset_name***. This page is similar to the Select options for importing an asset page on the asset import and update wizards.

Asset name

Specifies a logical name for the asset. An asset name must be unique within a cell and cannot contain an unsupported character.

An asset name cannot begin with a period (.), cannot contain leading or trailing spaces, and cannot contain any of the following characters:

Table 43. Characters that you cannot use in a name. The product does not support these characters in a name.

Unsupported characters		
/ forward slash	\$ dollar sign	' single quote mark
\ backslash	= equal sign	" double quote mark
* asterisk	% percent sign	vertical bar
, comma	+ plus sign	< left angle bracket
: colon	@ at sign	> right angle bracket
; semi-colon	# hash mark	& ampersand (and sign)
? question mark]]> No specific name exists for this character combination	

This **Asset name** field is the same as the **Name** setting on an Assets page.

Data type String

Asset description

Specifies a description for the asset.

Asset binaries destination URL

Specifies the directory to which the product imports the asset file.

Data type String
Units Full path name

Asset type aspects

Specifies the type of asset content. Examples of asset type include Java archive (JAR) files, shared libraries, enterprise application archive (EAR) files, and enterprise bundle archive (EBA) files.

The asset type suggests the content of the asset. For example an asset packaged as a JAR file might contain a web module, portlet and web service, and an asset packaged as an EBA file contains an OSGi application.

This setting is read-only. You cannot edit this setting.

Data type String
Units File type
Default none

File permissions

Specifies access permissions for asset binaries that the product expands to the asset binaries destination URL.

You can specify file permissions in the text field. You can also set some of the commonly used file permissions by selecting them from the list. List selections overwrite file permissions set in the text field.

You can set one or more of the following file permission strings in the list. Selecting multiple options combines the file permission strings.

Table 44. File permission string sets for list options. Select a list option or specify a file permission string in the text field.

Multiple-selection list option	File permission string set
Allow all files to be read but not written to	.*=755
Allow executables to execute	.*\.dll=755#.*\.so=755#.*\.a=755#.*\.sl=755
Allow HTML and image files to be read by everyone	.*\.htm=755#.*\.html=755#.*\.gif=755#.*\.jpg=755

Instead of using the multiple-selection list to specify file permissions, you can specify a file permission string in the text field. File permissions use a string that has the following format:

`file_name_pattern=permission#file_name_pattern=permission`

where `file_name_pattern` is a regular expression file name filter (for example, `.*\.jsp` for all JSP files), `permission` provides the file access control lists (ACLs), and `#` is the separator between multiple entries of `file_name_pattern` and `permission`. If `#` is a character in a `file_name_pattern` string, use `\#` instead.

If multiple file name patterns and file permissions in the string match a uniform resource identifier (URI) within the asset, then the product uses the most stringent applicable file permission for the file. For example, if the file permission string is `.*\.jsp=775#a.*\.jsp=754`, then the `abc.jsp` file has file permission 754.

Tip: Using regular expressions for file matching pattern compares an entire string URI against the specified file permission pattern. You must provide more precise matching patterns using regular expressions as defined by Java programming API. For example, suppose the product processes the following directory and file URIs during a file permission operation:

Table 45. Example URIs for file permission operations. Results are shown following this table.

1	/opt/WebSphere/profiles/AppSrv01/installedApps/MyCell/MyApp.ear/MyWarModule.war
2	/opt/WebSphere/profiles/AppSrv01/installedApps/MyCell/MyApp.ear/MyWarModule.war/MyJsp.jsp
3	/opt/WebSphere/profiles/AppSrv01/installedApps/MyCell/MyApp.ear/MyWarModule.war/META-INF/MANIFEST.MF
4	/opt/WebSphere/profiles/AppSrv01/installedApps/MyCell/MyApp.ear/MyWarModule.war/WEB-INF/classes/MyClass.class
5	/opt/WebSphere/profiles/AppSrv01/installedApps/MyCell/MyApp.ear/MyWarModule.war/mydir/MyClass2.class
6	/opt/WebSphere/profiles/AppSrv01/installedApps/MyCell/MyApp.ear/MyWarModule.war/META-INF

The file pattern matching results are:

- `MyWarModule.war` does not match any of the URIs
- `.*MyWarModule.war.*` matches all URIs
- `.*MyWarModule.war$` matches only URI 1
- `.*\.jsp=755` matches only URI 2
- `.*META-INF.*` matches URIs 3 and 6
- `.*MyWarModule.war/.*\.class` matches URIs 4 and 5

If you specify a directory name pattern for **File permissions**, then the directory permission is set based on the value specified. Otherwise, the **File permissions** value set on the directory is the same as its parent. For example, suppose you have the following file and directory structure:

`/opt/WebSphere/profiles/AppSrv01/installedApps/MyCell/MyApp.ear/MyWarModule.war/MyJsp.jsp`

and you specify the following file pattern string:

`.*MyApp.ear$=755#.*\.jsp=644`

The file pattern matching results are:

- Directory MyApp.ear is set to 755
- Directory MyWarModule.war is set to 755
- Directory MyWarModule.war is set to 755

Important: Regardless of the operation system, always use a forward slash (/) as a file path separator in file patterns.

Access permissions specified here are at the asset level. You can also specify access permissions for asset binaries in the node-level configuration. The node-level file permissions specify the maximum (most lenient) permissions that can be given to asset binaries. Access permissions specified here at asset level can only be the same as or more restrictive than those specified at the node level.

Data type String

Current asset relationships

Specifies the assets to which this asset is related.

To add or remove a relationship, use the Manage relationships page:

1. Click **Manage Relationships** to access the Manage relationships page. The **Selected** list on the right lists the current asset relationships.
2. To add a relationship, select an asset in the **Available** list on the left and click >>.
3. To remove a relationship, select an asset in the **Selected** list on the right and click <<.
4. Click **OK**.

Data type String

Default none

Validate asset

Specifies whether the product examines the asset references specified during asset importing or updating and, if validation is enabled, warns you of incorrect references or fails the operation.

An asset typically refers to resources using data sources for container-managed persistence (CMP) beans or using resource references or resource environment references defined in deployment descriptors. The validation checks whether the resource referred to by the asset is defined in the scope of the deployment target of that asset.

Select true (enable the check box) for resource validation and to stop operations that fail as a result of incorrect resource references. Select false (empty check box) for no resource validation.

Data type String

Default false (empty check box)

EBA Dependencies

For an enterprise bundle archive (EBA) asset, displays the current bundle download status for all bundles in the asset. This item is only displayed if your asset is an EBA asset, which means that it contains an OSGi application.

You cannot update an EBA asset until bundle downloads are complete from any previous update, and until the business-level application that uses the asset has picked up the previous updates by being restarted. Before you try and update bundle versions, you can use the EBA dependency information to check the bundle download status of the asset. The status displayed is one of the following values:

- Bundles downloading...

- Bundle downloads are complete.
- No bundles downloads are required.

Note: In addition to the information given here, you can also check the bundle download status indirectly, by checking the status of the associated EBA composition unit as described in Checking and updating the EBA asset version used by a business-level application.

If bundle downloads for the asset are complete, or no bundle downloads are required, you can update the asset using either of the methods described in Maintaining bundle versions for an EBA asset.

If bundle downloads for the asset are complete, and a new version of the EBA asset is available, restart the business-level application to bring the EBA composition unit up-to-date and to run the newer configuration.

Managing assets

After application binary files are imported and registered with the product management domain as assets, you can view, update and export those assets.

Before you begin

Import one or more assets. The name of each imported assets is shown on the list of assets on the administrative console Assets page.

About this task

You can view the contents of assets, update assets, remove assets from the product management domain, or export copies of assets to a target location. This topic describes how to perform these asset management operations from the administrative console Assets page. Alternatively, you can use programming or the wsadmin tool.

Procedure

- View or edit asset settings.
 1. Go to the administrative console Assets page.
Click **Applications > Application Types > Assets**.
 2. Click the asset name in the list of assets. The Asset settings page displays the values that are specified for the asset.
 3. Optional: Change the asset settings as needed and click **OK** to save the changes.
- Remove one or more assets from the product management domain.
- Update the contents of an asset.
- Export an asset to a target location.

What to do next

Create a business level application and add the asset to the business-level application.

Asset collection

Use this page to view a list of assets in the asset repository and to manage those assets. After importing an asset, you can add the asset to a business-level application.

Assets include Java archive (JAR) and compressed files that are used by applications installed on a server.

To view this administrative console page, click **Applications > Application Types > Assets**.

To view the values specified for an asset, click the asset name in the list. The displayed asset settings page shows the values specified. On the settings page, you can change existing asset values.

To manage an asset, enable the **Select** check box beside the asset name in the list and click a button:

Table 46. Button descriptions. Use the buttons to manage assets.

Button	Resulting action
Import	Opens a wizard that helps you add an asset to the asset repository.
Delete	Removes the asset from the asset repository and deletes the asset binaries from the file system of all nodes where the assets are installed. On single-server installations, deletion occurs after the configuration is saved. On multiple-server installations, deletion occurs after the configuration is saved and synchronized with the nodes.
Update	Opens a wizard that helps you update asset files. You can replace a file or module that exists on the server with a file or module that has the same name. Or you can add a new file or module, provided the new file or module does not have the same name as an asset that already exists on the server.
Export	Accesses the Export asset page, which you use to export an asset to a file at a location of your choice. Use the Export action to back up an asset.

Name

Specifies the name of the asset. Asset names must be unique within a cell and cannot contain an unsupported character.

Description

Specifies a description for the asset.

Updating assets

You can use the Update asset wizard to update classes, composites, wsdl, xsd, and definitions.xml files in an asset.

Before you begin

Import one or more assets. The file name of each deployable object in the imported assets is shown on the list of assets on the administrative console Assets page.

About this task

You can update all or part of the contents of assets that are in the product management domain. Complete the steps in the Procedure to update an asset using the administrative console Update asset wizard. Alternatively, you can update assets using programming or the wsadmin tool.

The following update limitations exist if the asset you are updating is a Service Component Architecture (SCA) asset:

- You cannot delete a composite file that a composition unit is using. If a delete is attempted, a warning message is sent to the Update asset log.
- You cannot update an sca-contribution.xml file.
 - SCA cannot detect deployable composites that are either added or deleted. Therefore, during deployment of a new composition unit, you do not see the new deployable composite in the deployables option list.

- SCA cannot detect dependencies that are added/removed during the Update asset process.
 - If a new import package is added and if a class in an existing composition unit is updated to require this new package, then the Update asset wizard fails with a `ClassNotFoundException`. Deployment of any new composition units from the updated asset are successful as the dependencies are detected during deployment operation.
 - If a new export package/namespace is added, then it has no affect on the existing composition unit and the Update asset wizard completes successfully.
- Because the Update asset wizard uses the new composite definition file provided in the asset for the existing composition unit, the following post deployment related changes to the composite configuration are not saved.
 - Binding resources: If you want to save this information, export all the data to the composite definition file in the new asset before you do the update.
 - Component reference target URIs: If you want to save this information, export all the data to the composite definition file in the new asset before you do the update.
 - Component properties: If you want to save this information, export all the data to the composite definition file in the new asset before you do the update.
 - HTTP Endpoint URL information: You need to reconfigure this information after the Update asset wizard finishes.
- For web services policy set attachments, during Update asset processing:
 - If there is a policy set specified for an endpoint in the updated composite definition file, SCA checks to see if a policy set has already been attached to that endpoint in the deployed composition unit. If an attachment already exists for that endpoint, the attachment is removed, and the policy set listed in the new composite file for that endpoint is attached. In this situation, if you have made any post deployment policy set configuration changes, these changes are lost.
 - If there is no policy set defined for an endpoint in the update composite definition file, then any existing attachments to that endpoint are removed.

Policy set bindings follow these same rules.
- For RunAs and RoleToUser mapping definitions, during Update asset processing:
 - For `implementation.java`, `implementation.spring` and `implementation.osgiapp`, any new roles defined in the `definition.xml` file in the asset are picked up and users can be mapped to these roles using either the `editCompositionUnit` command or the administrative console. Any existing role mappings for the original roles are preserved.
 - For `implementation.jee`, the `runAs` and `RoleToUser` mappings are defined in the JEE application instead of in the SCA asset or SCA composition unit. Therefore, SCA does not do anything with these mappings during Update asset processing.
- The user defined virtual host that hosts web content for `binding.ws`, `binding.atom`, `binding.http` with `wireformat.jsonrpc` and `implementation.widget` is not supported. A virtual host mapping of `default_host` is used during Update asset processing.

Procedure

1. Go to the Update asset wizard.
 - a. Click **Applications > Application Types > Assets** to access the Assets page.
 - b. Select the check box beside the asset that you want to update.
 - c. Click **Update**.
2. On the Update asset page, specify whether you want replace an entire asset or update its contents and, as needed, the replacement file or module.
 - a. Select an update option.

You can update asset contents by adding, deleting, or updating a single file or module in the asset, or by merging multiple files or modules. Update options include the following:

 - Replace entire asset

- Replace specific asset contents
- Add module or file to asset
- Remove file or module from asset
- Merge asset contents

The online help for the Update asset page describes the options.

- b. If you are updating specific asset contents or removing a file or module, specify the path beginning with the asset archive file.
 For **Specify the path beginning with the asset archive file**, specify a relative path to the file that starts from the root of the asset file. For example, if the file is located at `com/company/greeting.class` in module `hello.jar`, specify a relative path of `hello.jar/com/company/greeting.class`.
 - c. If you are updating the entire asset, updating an asset file or module, or merging asset contents, specify the full path name of the new file or module.
 - d. Click **Next**.
3. On the Select options for updating an asset page, specify asset settings and click **Next**.
 The online help for the Select options for importing an asset page describes the settings.
 4. On the Summary page, click **Finish**.

Results

If you update an asset packaged as a library JAR file that is not a Java Platform, Enterprise Edition (Java EE) archive, then the product automatically distributes the updated asset to all of the composition units that use the asset.

However, if you update a Java EE asset, then the product does not automatically distribute the updated Java EE archive to composition units created from that asset, which are Java EE applications. You must select every Java EE application created from that asset and use the **Update** button to update the Java EE application individually by specifying the update contents.

What to do next

Create a business-level application and add the asset to the business-level application.

Update asset settings

Use this page to select whether you want replace an entire asset or update its contents. You can update asset contents by adding, deleting, or updating a single file or module in the asset, or by merging multiple files or modules into an asset. Updating an asset registers the updated files with the product management domain.

To view this administrative console page, click **Applications > Application Types > Assets**, select the asset to update, and then click **Update**.

The product manages the contents of a registered asset as a single entity. The contents of a registered asset must be accessible to application servers, web servers and other runtime environments that use the asset.

When you replace an asset or update an asset by adding a file or module, asset files typically are uploaded from a client workstation running the browser to the server machine running the administrative console, where they are registered. In such cases, use the web browser running the administrative console to select files to upload to the server machine.

The specified asset that you are installing must be one of the following supported assets:

- A single file, such as an enterprise bean (EJB) file
- An archive of files, such as a Java archive (JAR) or a compressed .zip file

- An archive of archives, such as an enterprise archive (EAR) or shared library file

Replace entire asset:

Under **Select the type of update to perform**, specifies to replace the entire asset installed on the server with a new (updated) asset.

After selecting this option, specify whether the asset is on a local or remote file system and the full path name of the asset. The path provides the location of the updated asset before installation.

Use **Local file system** if the browser and asset files are on the same machine (whether or not the server is on that machine, too).

Use **Remote file system** if the asset file resides on any node in the current cell context. Only supported assets are shown during the browsing. Also use **Remote file system** to specify an asset file that is already residing on the machine running the application server. For example, the field value might be *profile_root/installableApps/my_bean.ejb*. After the asset file is transferred, the **Remote file system** value shows the path of the temporary location on the server.

Replace specific asset contents:

Under **Select the type of update to perform**, specifies to replace a file or module of the asset installed on the server.

After selecting this option, do the following:

1. For **Specify the path beginning with the asset archive file**, specify a relative path to the file that starts from the root of the asset file. For example, if the file is located at *com/company/greeting.class* in module *hello.jar*, specify a relative path of *hello.jar/com/company/greeting.class*.
2. Specify whether the asset is on a local or remote file system and the full path name of the asset. The path provides the location of the updated asset before installation.
3. Click **Next**.

The **Replace entire asset** description describes options for specifying the full path name of an asset or file to add using **Local file system** and **Remote file system** options.

Add a module or file to an asset:

Under **Select the type of update to perform**, specifies to add a file to the asset installed on the server.

After selecting this option, do the following:

1. For **Specify the path beginning with the asset archive file**, specify a relative path to the file that starts from the root of the asset file. For example, if the file is located at *com/company/greeting.class* in module *hello.jar*, specify a relative path of *hello.jar/com/company/greeting.class*.
2. Specify whether the asset is on a local or remote file system and the full path name of the asset. The path provides the location of the updated asset before installation.

The **Replace entire asset** description describes options for specifying the full path name of an asset or file to add using **Local file system** and **Remote file system** options.

Remove a file or module from an asset:

Under **Select the type of update to perform**, specifies to remove a file or module from the asset installed on the server.

After selecting this option, do the following:

1. For **Specify the path beginning with the asset archive file**, specify a relative path to the file to be removed that starts from the root of the asset file. For example, if the file is located at `com/company/greeting.class` in module `hello.jar`, specify a relative path of `hello.jar/com/company/greeting.class`.
2. Click **Next**.

Merge asset contents:

Under **Select the type of update to perform**, specifies to compare the new file or module with the file or module of the asset installed on the server. If the file or module exists, it is replaced. Otherwise, it is added to the installed asset.

After selecting this option, specify whether the new file or module is on a local or remote file system and the full path name of the file or module. The path provides the location of the updated asset before installation.

The **Replace entire asset** description describes options for specifying the full path name of a file or module to merge using **Local file system** and **Remote file system** options.

Update associated composition unit:

Specifies whether to update the composition units that are associated with an enterprise (Java EE) asset. This option applies to enterprise assets only.

The default value is `NONE`. Specify `ALL` to update all of the composition units that are associated with the enterprise asset.

Deleting assets

You can remove application binary files that are registered as assets from the product management domain.

Before you begin

Import one or more assets. The name of each imported asset is shown on the list of assets on the administrative console Assets page.

About this task

You can remove assets from the product management domain, provided the asset does not have an existing composition unit. If an asset has one or more composition units defined in the management domain, then you cannot delete that asset until those composition units are removed.

This topic describes how to delete assets using the administrative console. Alternatively, you can use programming or the `wsadmin` tool.

Procedure

1. Go to the Delete asset page.
 - a. Click **Applications > Application Types > Assets** to access the Assets page.
 - b. Select the check box beside the asset that you want to delete.
 - c. Click **Delete**.
2. On the Delete asset page, click **OK** to confirm that you want the specified asset removed from the product management domain.
Click **Cancel** to return to the Assets page and not delete the asset.

Results

The product deletes the asset from the product management domain.

In a multiple-server environment, the product deletes the asset binary files from the target node machine after node synchronization.

What to do next

On the Assets page, verify that the deleted asset is no longer in the list of imported assets.

Exporting assets

After application binary files are imported and registered with the product management domain as assets, you can export those assets.

Before you begin

Import one or more assets. The file name of each deployable object in the imported assets is shown on the list of assets on the administrative console Assets page.

About this task

You can export copies of assets to a target location. Exporting stores application binary files, enabling you to back up the files or edit them. The file resulting from exporting an asset contains configuration information for the asset.

This topic describes how to export an asset from the administrative console Assets page. Alternatively, you can use programming or the wsadmin tool.

Procedure

1. Go to the Export asset page.
 - a. Click **Applications > Application Types > Assets** to access the Assets page.
 - b. Select the check box beside the asset that you want to export.
 - c. Click **Export**.
2. On the Export asset page, click the asset name or identifier.

To cancel the export operation and return to the Assets page, click **Back**.
3. Specify the target location for the asset file.

What to do next

Examine the target file to verify that the asset exported correctly. You can later edit this file and import the edited asset.

Creating business-level applications

You can create an empty business-level application and then add assets, shared libraries, business-level applications, and other artifacts as composition units to the empty business-level application.

Before you begin

Configure each target application server as needed. You must deploy a business-level application to a Version 7.0 server or cluster.

Optionally, determine what assets or other files that you want to add to your business-level application and whether your application files can run on your deployment targets.

About this task

You can create business-level applications using the administrative console, programming, or the wsadmin tool.

Procedure

1. Select a way to create your business level application.

Table 47. Ways to create business level applications. You can create business-level applications using the administrative console, programming, or wsadmin.

Option	Method
Administrative console business-level application creation wizard See "Creating business-level applications with the console."	Click Applications > New application > New Business-level Application and follow instructions in the wizard.
Administrative console Java Platform, Enterprise Edition (Java EE) application installation wizard See "Installing enterprise application files with the console" on page 213.	Click Applications > New application > New Enterprise Application and follow instructions in the wizard. The product creates a new business-level application with the enterprise application that you install or makes the enterprise application a composition unit of an existing business-level application. See the Business-level application name setting on the Select installation options wizard page.

2. Create your business-level application using the administrative console, programming or wsadmin.
3. Save the changes to your administrative configuration.
When saving the configuration, synchronize the configuration with the nodes where the application is expected to run.

Results

The name of the application is shown in the list on the Business-level applications page.

What to do next

After you create a business-level application, you can do the following to add composition units to it:

1. Import any assets needed by your business-level application.
2. Add assets, shared libraries, or other business-level applications as composition units.
3. Save the changes to your administrative configuration.
4. Start the business-level application.

If the application does not run as desired, edit the application configuration, then save and run it again.

Creating business-level applications with the console

You can create an empty business-level application and then add assets or business-level applications as composition units to the empty business-level application.

Before you begin

Before you create a business-level application, decide upon an application name. Optionally, determine which assets, shared libraries, or business-level applications that the new business-level application needs.

About this task

This topic describes how to create an empty business-level application and then add assets as composition units to the application using the administrative console. Alternatively, you can use programming or the wsadmin tool.

You can add an asset or shared library composition unit to multiple business-level applications. However, each composition unit for the same asset must have a unique composition unit name. You can add a business-level application composition unit to more than one business-level application.

Procedure

1. Create an empty business-level application.
 - a. Click **Applications > New application > New Business Level Application**.
 - b. On the New business-level application page, specify a unique name for the application and a description, and then click **OK**.
 - c. On the business-level application settings page, click **Save**.

The name and description are shown in the list of applications on the Business-level applications page. Because the application is empty, its status is `Unavailable`.
2. Optional: Add one or more assets, non-Java EE shared libraries, or business-level applications to a business-level application. The product adds these assets as composition units of your business-level application.
 - If the asset that you want to add to your business-level application is a Java Platform, Enterprise Edition (Java EE) application or module that is not yet deployed, see step 3.
 - If the asset is a Java EE shared library, see step 4.
 - If the asset is an enterprise bundle archive (EBA) asset, see Adding an EBA asset to a business-level application using the administrative console.
 - a. Import the assets or create the business-level applications that you want to add to the business-level application.
 - b. Go to the business-level application settings page.

Click **Applications > Application Types > Business-level applications > *application_name***.
 - c. On the business-level application settings page, specify the type of composition unit to add.
 - To add an asset, under **Deployed assets**, click **Add > Add Asset**.
 - To add a shared library, under **Deployed assets**, click **Add > Add Shared Library**.
 - To add a business-level application, under **Business-level applications**, click **Add**.
 - d. On the Add page, select a unit from the list of available units, and then click **Continue**.

If you are adding one or more deployable unit assets and you have multiple imported assets available, you can select more than one deployable unit.
 - e. On the Set options page, change the composition unit settings as needed, and then click **Next**.

This page is not shown when you add a Java EE asset as a shared library or if you have multiple deployable unit assets. If the application installation or update wizard displays and you want to add a Java EE asset as a shared library, see step 4.
 - f. On the Map composition unit to a target page, change the deployment target as needed, and then click **Next**.

This page is not shown when you add a business-level application.

g. If you are adding one or more deployable unit assets, specify composition unit relationship options. See “Relationship options settings” on page 414.

h. On the Summary page, click **Finish**. Several messages are displayed, indicating whether the product adds the unit to the business-level application successfully. A message having the format `Completed res=[WebSphere:cuname=unit_name,cuedition=version]` indicates that the addition is successful. Click **Manage application**.

If the product adds the unit successfully, the name of the unit is shown on the list of composition units on the Adding composition unit to the business-level application page.

If the unit addition is not successful, read the messages and try adding the unit again. Correct the problems noted in the messages.

i. On the Adding composition unit to the business-level application page, click **Save**.

The product creates composition units for the asset, shared library, or business-level application. The unit names are shown in lists of composition units on the settings page of your business-level application. To view the settings page, click **Applications > Application Types > Business-level applications > *your_application_name***.

3. Optional: Install a Java EE application or module, and add it as a composition unit to your business-level application.

When installing an enterprise archive (EAR) file or a stand-alone Java EE module using the application installation wizard, you can specify a business-level application to which to add the EAR file or module. You can also specify relationships to any shared libraries that your Java EE application or module uses. The product creates composition units that represent those relationships.

a. Click **Applications > New application > New Enterprise Application**.

b. On the first Preparing for the application installation page, specify the Java EE application or module to install and click **Next**.

c. On the second Preparing for the application installation page, select **Detailed - Show all installation options and parameters**, specify whether to generate default bindings and mappings as needed for the application or module, and click **Next**.

d. On the Select installation options page of the wizard, select your business-level application for **Business-level application name** and click **Next**. The product creates a composition unit that has the same name as the Java EE application or module and adds the unit to your business-level application.

If you do not specify a value for **Business-level application name**, then the product creates a default business-level application that has the same name as the Java EE application that you are installing. The product does not add the Java EE application as a composition unit to the business-level application that you created in step 1.

e. Optional: On the Map shared library relationship page of the wizard, specify relationship identifiers and composition unit names for shared libraries that modules in your Java EE application use. The product creates a composition unit for each shared library relationship in your business-level application.

You can map shared library relationships when installing your Java EE application or module or, after installation, return to the Map shared library relationship page and specify shared library relationships. See step 4.

f. Complete the other application installation wizard options as needed to install the Java EE application or module.

The product creates composition units for the application, module, or shared library relationships. The unit names are shown in lists of composition units on the settings page of your business-level application. To view the settings page, click **Applications > Application Types > Business-level applications > *your_application_name***.

4. Optional: After installation of a Java EE application or module, you can specify composition units for relationships to shared libraries that are used by your business-level application. Specify relationships to shared libraries on the Map shared library relationship page of the application installation or update wizard.
 - a. If you have not done so already, import a Java EE asset such as an enterprise bean (EJB) or web module (WAR) that uses a shared library file.

If the product displays `javaarchive` for **Asset type aspects** on the asset settings page, continue to step 4b.

If the product does not display `javaarchive` for **Asset type aspects** on the asset settings page, then the asset is not a Java EE asset. Use step 2 to add a shared library to your business-level application.
 - b. Go to a settings page for your business-level application.

Click **Applications > Application Types > Business-level applications > *your_application_name***.
 - c. Under **Deployed assets**, click **Add > Add Shared Library**.
 - d. On the Add composition unit page, select the Java EE asset that you imported and then click **Continue**.

The Java EE application installation or update wizard displays. Select the Java EE application or module that uses the asset, and complete the steps in the wizard.
 - e. On the Select installation options page of the wizard, select your business-level application for **Business-level application name**.
 - f. On the Map shared library relationship page of the wizard, specify a relationship identifier and composition unit name for the asset.
 - g. Complete the other wizard options as needed.

The product creates a composition unit for the shared library relationship. The unit name is shown in the list of deployed asset composition units on the settings page of your business-level application.

Results

The name of your business-level application is shown on the Business-level applications page in the list of applications.

What to do next

After you create the application, save the changes to your configuration and start the application as needed.

Business-level application collection

Use this page to view and manage business-level applications.

To view this administrative console page, click **Applications > Application Types > Business-level applications**.

To view the values specified for an application configuration, click the application name in the list. The displayed application settings page shows the values specified. On the settings page, you can change existing configuration values and link to additional console pages that assist you in configuring the application.

To manage a business-level application, enable the **Select** check box beside the application name in the list and click a button:

Table 48. Button descriptions. Use the buttons to manage business-level applications.

Button	Resulting action
Start	Attempts to run the application. After the application starts successfully, the state of the application changes to <i>Started</i> if the application starts on all deployment targets, else the state changes to <i>Partial Start</i> .
Stop	Attempts to stop the processing of the application. After the application stops successfully, the state of the application changes to <i>Stopped</i> if the application stops on all deployment targets, else the state changes to <i>Partial Stop</i> .
New	Opens a wizard that helps you add assets, shared libraries, or business-level applications as composition units to your application.
Delete	Deletes the application from the product configuration repository and deletes the application binaries from the file system of all nodes where the application modules are installed. On single-server installations, deletion occurs after the configuration is saved. On multiple-server installations, deletion occurs after the configuration is saved and synchronized with the nodes.

Name:

Specifies the name of the business-level application. Application names must be unique within a cell and cannot contain an unsupported character.







Description:

Specifies a description for the business-level application.

Status:

Indicates whether the application deployed on the application server is started, stopped, or unknown.

Table 49. Application status. The status indicates whether the application is running.

	Started	Application is running.
	Partial start	Application is in the process of changing from a <i>Stopped</i> state to a <i>Started</i> state. Application is starting to run but is not fully running yet. Or, it cannot fully start because a server mapped to one or more application modules is stopped.
	Stopped	Application is not running.
	Partial stop	Application is in the process of changing from a <i>Started</i> state to a <i>Stopped</i> state. Application has not stopped running yet.
	Unknown	Status cannot be determined.
	Pending	Status is temporarily unknown pending an event that a user did not initiate, such as pending an asynchronous call.
	Not applicable	Application does not provide information as to whether it is running.

The status of an application on a web server is always **Unknown**.

New business-level application settings

Use this page to name and describe a new business-level application.

To view this administrative console page, click **Applications > New application > New Business-level Application**.

Name:

Specifies a logical name for the business-level application. An application name must be unique within a cell and cannot contain an unsupported character.

An application name cannot begin with a period (.), cannot contain leading or trailing spaces, and cannot contain any of the following characters:

Table 50. Characters that you cannot use in a name. The product does not support these characters in a name.

Unsupported characters		
/ forward slash	\$ dollar sign	' single quote mark
\ backslash	= equal sign	" double quote mark
* asterisk	% percent sign	vertical bar
, comma	+ plus sign	< left angle bracket
: colon	@ at sign	> right angle bracket
; semi-colon	# hash mark	& ampersand (and sign)
? question mark]]> No specific name exists for this character combination	

Data type String

Description:

Specifies a description for the application.

This field is the same as the **Description** setting on a Business-level applications page.

Shared library relationship and mapping settings

Use the Shared library relationship and Shared library relationship mapping pages to specify relationship identifiers and composition unit names for shared libraries that modules in your enterprise application reference. When installing your enterprise application, the product creates a composition unit for each shared library relationship in the business-level application that you specified on the Select installation options page of the application installation wizard.

To view this console page in a wizard, click **Applications > Install new application > New Enterprise Application > application_path > Next > Detailed - Show all installation options and parameters > Next > application_name > Step: Map shared library relationships.**

After installation, click **Applications > Application Types > WebSphere enterprise applications > Shared library relationships.**

To map library files used in a business-level application to an application or web module, use the Shared library relationship mapping page:

1. Click **Reference shared libraries.**
2. Note the application or module in **Map libraries to the application or module listed.** You are associating library files with that application or module.
3. From the **Available** list, select one or more libraries that the application or module uses.
4. Click >> to add them to the **Selected** list.
5. To remove an association, select one or more libraries in the **Selected** list and click <<.
6. Click **OK.**

Module:

Specifies the name of the module associated with the shared libraries.

URI:

Specifies the location of the module relative to the root of the application EAR file.

Relationship identifiers:

Specifies an identifier for a module shared library relationship. The product assigns an identifier to the composition unit that it creates for the shared library relationship in the business-level application.

Composition unit names:

Specifies a composition unit name for the shared library relationship. The product uses this value to name the composition unit that it creates for the shared library relationship in the business-level application that you specified on the Select installation options page of this wizard.

This setting is only in the application installation and update wizards.

Match target:

Specifies whether the product maps the composition unit for the shared library relationship to the same deployment target as the business-level application.

Note: If you later change the deployment target of the business-level application or its modules, you must manually update the shared library target to match the target of the application and modules. The targets of shared library composition units are not automatically updated. Not updating the target of the shared library composition unit might cause `java.lang.ClassNotFoundException` errors and prevent the application or its modules from starting. To prevent these error conditions, also ensure that shared libraries upon which other modules or applications depend have a lower starting weight than dependent applications and modules.

Add composition unit settings

Use this page to specify options for the composition unit to be added to the business-level application. The product assigns a default value for an option when you do not specify a value.

To view this administrative console page, click **Applications > Application Types > Business-level applications > *business-level_application_name* > Add > Add unit_type**.

Name:

Specifies the name of the composition unit to be added to the business-level application.

The table lists available composition units. Select a unit from this list.

Description:

Specifies a description for the composition unit.

Add asset settings

Use this page to add one or more assets to a business-level application.

To view this administrative console page, click **Applications > Application Types > Business-level applications > *application_name* > Add > Add Asset**.

Deployable units:

Specifies the imported assets available for use in a business-level application. The list of deployable units includes only imported assets, and not shared libraries or business-level applications.

From this list, select one or more deployable units to add as composition units to your business-level application.

Set options settings

Use this page to specify options for the composition unit to be added to the business-level application. The product supplies default values for the options if you do not specify a value.

To view this administrative console page, click **Applications > Application Types > Business-level applications > application_name**. On the business-level application settings page, specify the type of composition unit to add:

- To add an asset, under **Deployed assets**, click **Add > Add Asset**.
- To add a shared library, under **Deployed assets**, click **Add > Add Shared Library**.
- To add a business-level application, under **Business-level applications**, click **Add**.

Backing identifier:

Specifies a unique identifier for a composition unit that is registered in the application management domain.

The identifier has the format: `WebSphere:unit_typename=unit_name,unit_typeversion=version_number`. For example, for the `MyApp.jar` asset, the backing identifier might be `WebSphere:assetname=MyApp.jar`.

Data type String
Units Composition unit identifier

Name:

Specifies the name of the composition unit.

For example, for the `MyApp.jar` asset, the name might be `MyApp.jar`.

A unit name cannot begin with a period (.), cannot contain leading or trailing spaces, and cannot contain any of the following characters:

Table 51. Characters that you cannot use in a name. The product does not support these characters in a name.

Unsupported characters		
/ forward slash	\$ dollar sign	' single quote mark
\ backslash	= equal sign	" double quote mark
* asterisk	% percent sign	vertical bar
, comma	+ plus sign	< left angle bracket
: colon	@ at sign	> right angle bracket
; semi-colon	# hash mark	& ampersand (and sign)
? question mark]]> No specific name exists for this character combination	

Data type String

Description:

Specifies a description for the composition unit.

Starting weight:

Specifies the order in which composition units are started when the server starts. The starting weight is like the startup order. The composition unit with the lowest starting weight is started first.

The value that you set for **Starting weight** determines the importance or weight of a composition unit within the business-level application. For example, for the most important composition unit within a business-level application, specify 1 for **Starting weight**. For the next most important composition unit within the business-level application, specify 2 for **Starting weight**, and so on.

Data type	Integer
Default	1
Range	0 to 2147483647

Start composition unit upon distribution:

Specifies whether to start the composition unit after the product distributes the composition unit to other locations.

The default is not to start the composition unit.

Data type	Boolean
Default	false

Restart behavior on update:

Specifies whether the product restarts deployment targets after updates to the composition unit.

Usually, a composition unit is mapped to one or more deployment targets. This setting determines whether the product restarts those targets after editing the composition unit.

Table 52. Restart behavior on update options. Depending on your selection, the product restarts all target nodes, the nodes controlled by sync plug-ins, or no nodes.

Option	Description
ALL	The product restarts each target node of the composition unit after editing the composition unit.
DEFAULT	The product restarts the nodes controlled by the sync plug-ins after editing the composition unit.
NONE	The product does not restart nodes after editing the composition unit.

Map target settings

Use this page to map a composition unit to a deployment target. The product assigns a default target when you do not specify a target.

To view this administrative console page, click **Applications > Application Types > Business-level applications > application_name > composition_unit_name > Modify Target**. The Map target page is similar to the Map composition unit to a target page in the add composition unit wizard.

On single-server products, a deployment target can be an application server or web server.

On multiple-server products, a deployment target can be an application server, cluster of application servers, or web server.

On this page, map a composition unit to one or more desired targets.

Current targets:

Specifies the existing deployment targets for the composition unit.

Available:

Lists the names of available deployment targets. This list is the same for every composition unit that is registered in the cell.

From this list, select only appropriate deployment targets for a composition unit.

If the unit calls a Version 8.x application programming interface (API) or uses a 8.x feature, then you must map the unit to a 8.x deployment target. If the unit supports Java Platform, Enterprise Edition (Java EE) 6, then you must map the unit to a 8.x deployment target.

If the unit calls a Version 7.x API, uses a 7.x feature, or supports Java EE 5, then you must map the unit to an 8.x or 7.x deployment target.

If the unit supports Java 2 Platform, Enterprise Edition (J2EE) 1.4, then you must map the unit to an 8.x, 7.x or 6.x deployment target. You can map units that call a 6.x API or use a 6.x feature to an 8.x, 7.x or 6.x deployment target.

To map a composition unit to a deployment target, select a target from the **Available** list and click >>. The target name is displayed in the **Selected** list.

Selected:

Lists the names of desired deployment targets.

When you click **OK**, the product maps the composition unit to the deployment targets in the **Selected** list.

To remove a deployment target from the **Selected** list, select the target and click <<.

Relationship options settings

Use this page to specify relationship options for deployable or composition units in an asset deployed as part of a business-level application. Specifying a relationship declares a dependency relationship that a deployable unit or composition unit has on another asset deployed as a shared library in the same business-level application.

To view this administrative console page, click **Applications > Application Types > Business-level applications > application_name > deployed_asset_name > Relationship options**. This help also pertains to wizard pages that are shown when you add multiple deployable or composition unit assets to a business-level application. These pages are shown for the **Define relationship with existing composition units** and **Options for creating new composition units to satisfy asset relationships** wizard steps.

A business-level application consists of composition units. When you add an asset to a business-level application, the product creates a composition unit for the asset. The composition unit name can be different from the name of the asset being deployed. The list of deployed assets shown for a business-level application consists of the composition unit names for the deployed assets. The relationships defined in this page are composition unit relationships. The deployable units listed for a composition unit are those you chose from the associated asset when adding the asset. Composition unit relationships are expressed as deployable unit dependencies on other composition units belonging to the same business-level application. Only a composition unit for an asset deployed as a shared library can be specified as a dependency. You can map each deployable unit to a target independently from the others. Modifying relationships in this page only affects the composition unit, not the associated asset.

To specify relationship options, select a deployable unit and click a button.

Button	Resulting action
Set Relationships	Displays a page through which you can add or change relationships for the deployable unit. Specify a relationship if a deployable unit depends on another asset deployed as a shared library in order to run. This button is on the Set relationship options page.
Enable Match Targets	If the deployable unit has a dependency relationship defined, click Enable Match Targets to map the related deployed assets to the same deployment targets as the dependent deployable unit.
Disable Match Targets	If the deployable unit has a dependency relationship defined, click Disable Match Targets if the related deployed assets do not need to be deployed to the same targets as the deployable unit.

Deployable unit name or composition unit name:

Specifies the name of the deployable unit or the composition unit of the selected deployed asset.

Relationship:

Specifies the composition unit names for all relationships defined for the associated deployable unit.

This setting is on the Set relationship options page.

By default, a deployable unit has no relationships. To add or change related composition units, do the following:

1. Select the deployable unit.
2. Click **Set Relationships**.
3. Select the composition units that the deployable unit requires by moving them from the **Available** list to the **Selected** list.
4. Click **OK**.

Match targets:

Indicates the match targets value selected for the associated deployable unit. The default value is true.

A match targets value of true maps the composition units listed under **Relationship** to the same deployment targets as the associated deployable unit. Typically, you must deploy related composition units to the same targets as the dependent deployable unit in order for the deployable unit to run.

A false value indicates that the related composition unit can map to deployment targets which are different from the deployment targets of the deployable unit.

To set the value to true, select the deployable unit and click **Enable Match Targets**. To set the value to false, select the deployable unit and click **Disable Match Targets**. To set this value, the deployable unit must have a related composition unit.

If you have multiple deployable units and your security configuration supports multiple domains, the deployable units must be in the same security domain.

Business-level application settings

Use this page to configure a business-level application.

To view this administrative console page, click **Applications > Application Types > Business-level applications > application_name**.

This page is the same as the Adding composition unit to the business-level application page.

Name

Specifies a logical name for the application. An application name must be unique within a cell and cannot contain an unsupported character.

An application name cannot begin with a period (.), cannot contain leading or trailing spaces, and cannot contain any of the following characters:

Table 53. Characters that you cannot use in a name. The product does not support these characters in a name.

Unsupported characters		
/ forward slash	\$ dollar sign	' single quote mark
\ backslash	= equal sign	" double quote mark
* asterisk	% percent sign	vertical bar
, comma	+ plus sign	< left angle bracket
: colon	@ at sign	> right angle bracket
; semi-colon	# hash mark	& ampersand (and sign)
? question mark]]> No specific name exists for this character combination	

Data type String

Description

Specifies a description for the business-level application.

Deployed assets

Specifies the asset and shared library composition units in the business-level application. A *composition unit* is a registered asset or shared library that has additional configuration information, which you specify when adding the asset to the application.

For each composition unit, the table provides a name, description, asset type, and the runtime status of the composition unit.

Table 54. Deployed assets button descriptions. Use the buttons to add or delete composition units.

Button	Resulting action
Add > Add Asset	For assets that contain Java Platform, Enterprise Edition (Java EE) applications or modules, opens the application installation wizard. On the Select installation options page of this wizard, you can specify a Business-level application name value that identifies the target business-level application. On the Map shared library relationships page, you can identify the shared library files that individual modules need to run and specify composition unit names for the module-shared library relationships. For non-Java EE assets, opens a wizard that helps you add an asset as a composition unit to your business-level application.
Add > Add Shared Library	Opens a wizard that helps you add a library file as a composition unit to your business-level application.
Delete	Deletes the composition unit from the product configuration repository and deletes the application binaries from the file system of all nodes where the application modules are installed. On multiple-server installations, deletion occurs after the configuration is saved and synchronized with the nodes.

Business-level applications

Specifies the business-level applications in this business-level application.

The table provides a name, description, and the runtime status of each contained business-level application.

Table 55. Business-level applications button descriptions. Use the buttons to add or delete composition units.

Button	Resulting action
Add	Opens a wizard that helps you add a business-level application to your business-level application.
Delete	Deletes the business-level application from the product configuration repository and deletes the application binaries from the file system of all nodes where the application modules are installed. On multiple-server installations, deletion occurs after the configuration is saved and synchronized with the nodes.

Composition unit settings

Use this page to view composition unit settings and to change the configuration properties of a composition unit. The specific settings that are available for configuration can vary, depending upon the contents of the composition unit. For example, there are additional configuration settings if the asset contained in the composition unit is an SCA composite, or an OSGi application.

To view this administrative console page, click **Applications > Application Types > Business-level applications > application_name > deployed_asset_name**. The deployed asset is a composition unit of the business-level application.

- “Settings that are common to all composition units”
- “Additional composition unit settings for SCA composites” on page 419
- “Additional composition unit settings for OSGi applications” on page 419

Settings that are common to all composition units

Name:

Specifies a logical name for the composition unit. You cannot change the name on this page.

Description:

Specifies a description for the composition unit.

Backing ID:

Specifies a unique identifier for a composition unit that is registered in the application management domain.

The identifier has the format `WebSphere:unittypename=unit_name`. For example, for the `MyApp.jar` asset, the backing identifier might be `WebSphere:assetname=MyApp.jar`.

You cannot change the identifier on this page.

Data type String
Units Configuration unit identifier

Starting weight:

Specifies the order in which composition units are started when the server starts. The starting weight is like the startup order. The composition unit with the lowest starting weight is started first.

The value that you set for **Starting weight** determines the importance or weight of a composition unit within the business level application. For example, for the most important composition unit within a business-level application, specify 1 for **Starting weight**. For the next most important composition unit within the business-level application, specify 2 for **Starting weight**, and so on.

Note: Assign composition units upon which other composition units depend a lower starting weight than the dependent composition units. If a composition unit is not started and running before its dependent composition units, `java.lang.ClassNotFoundException` errors might result when you attempt to start the application or its modules.

Data type	Integer
Default	1
Range	0 to 2147483647

Start on distribution:

Specifies whether to start the composition unit when the product distributes the composition unit to other locations.

The default is not to start the composition unit.

This setting applies to asset or shared library composition units. This setting does not apply when the composition unit is a business-level application.

Data type	Boolean
Default	false

Recycle behavior on update:

Specifies whether the product restarts the composition unit after the composition unit is updated.

The default is to restart the composition unit after partial updating of the composition unit.

This setting applies to asset or shared library composition units. This setting does not apply when the composition unit is a business-level application.

Table 56. Option descriptions. Specifies whether to restart an asset or shared library composition unit.

Option	Description
ALL	Restarts the composition unit after the entire composition unit is updated
DEFAULT	Restarts the composition unit after the part of the composition unit is updated
NONE	Does not restart the composition unit after the composition unit is updated

Target mapping:

Specifies the current targets for the composition unit.

To change the deployment targets, click **Modify targets** then select a different set of deployment targets from the list of available clusters and servers.

For SCA, you must specify only a single server or cluster as the target. Do not map an SCA composition unit to multiple servers or clusters.

Note: When you change the deployment target of composition units in a business-level application, the startup order changes to the same order in which you remap composition unit targets, even if the starting weight for all composition units is set to 1. To avoid `java.lang.ClassNotFoundException` errors when attempting to start the remapped composition units, remap targets for composition units in the same order as that used to add the composition units or, after remapping, check starting weights to ensure that composition units upon which other composition units depend are started first.

Additional composition unit settings for SCA composites

SCA composite components:

Specifies the component names and component implementations of SCA composites in the application.

Table 57. Column descriptions. Provides the name of each component and the name of the class or code implementing the component.

Column	Description
Component Name	Specifies the name of a component associated with the SCA composite.
Component Implementation	Specifies the name of the class or code implementing the component.

None indicates that the SCA composite does not have defined components.

SCA composite properties:

Specifies the names and values of SCA composite properties in the application.

Table 58. Column descriptions. Provides the name and value of SCA composite properties.

Column	Description
Property Name	Specifies the name of an SCA composite property.
Property Value	Specifies the value of the property.

None indicates that the SCA composite does not have defined name-value properties.

SCA composite wires:

Specifies the sources and targets of wires in the SCA composite.

Table 59. Column descriptions. Provides the source and target of wires.

Column	Description
Wire Source	Specifies the source of a wire in the SCA composite.
Wire Target	Specifies the target of the wire.

None indicates that the SCA composite does not have defined wires.

Additional composition unit settings for OSGi applications

OSGi application deployment status:

The deployment status shows whether updates are available for the EBA asset that is contained in the composition unit. If a new version of an EBA asset is available, and all bundle downloads for the asset are complete, you can update the EBA composition unit so that the business-level application uses the latest configuration. You do not have to update the composition unit every time you update the asset.

There are four distinct deployment statuses for an EBA composition unit:

Using latest OSGi application deployment.

The composition unit is running the latest configuration of the backing asset and any CBA extensions.

New OSGi application deployment not yet available because it requires bundles that are still downloading.

The backing asset is currently undergoing a bundle version update, or bundles are downloading for a CBA extension.

New OSGi application deployment available.

The backing asset is available at a newer configuration than the configuration that is currently running in this composition unit, or a CBA extension has been added or replaced.

New OSGi application deployment cannot be applied because bundle downloads have failed.

The last bundle version update for the backing asset or CBA extension did not succeed, and therefore the newer configuration is not yet available.

If the status is “New OSGi application deployment available”, the **Update to latest deployment ...** button is available. Click this button to bring the EBA composition unit up-to-date and run the updated business-level application. If any of the updates need configuration changes, a wizard prompts you to update the configuration information.

When you save the changes to the EBA composition unit, the associated business-level application is updated to use the new configuration. If the business-level application is running, the bundle and configuration updates are applied immediately. If possible (that is, depending on the nature of the updates) the system applies the updates without restarting the application. Updates that pull in new use bundles at run time prompt a full restart of the application. Updates that pull in new provision bundles might also prompt a full application restart.

Example: Creating a business-level application

You can add many different types of artifacts to business-level applications. For example, you can add Java Platform, Enterprise Edition (Java EE) applications or modules, Java archives (JAR files), data in compressed files, and other business-level applications.

About this task

An example of creating a simple business-level application follows. This example assumes that you have a compressed file, such as a compressed file, or other archive available on your computer or on a remote server that you can use to complete the example.

If you do not have a compressed file available, look in product directories. Installing the product samples adds several sample files to the `/samples` directory. You can use these sample files in a business-level application.

Procedure

1. Import assets.
 - a. Click **Applications > New application > New Asset** in the console navigation tree.
 - b. On the Upload asset page, specify the asset package to import and click **Next**.
For example, specify a compressed file such as a compressed file and click **Next**.

- c. On the Select options for importing an asset page, click **Next**.
- d. On the Summary page, click **Finish**.
- e. On the Adding asset to repository page, if messages show that the operation completed, click **Manage assets**.
- f. On the Assets page, click **Save**.

The file name displays in the list of assets.

2. Create an empty business-level application named MySampleBLA.
 - a. Click **Applications > New application > New Business Level Application**.
 - b. On the New business-level application page, specify a unique name such as MySampleBLA and a description, and then click **OK**.
 - c. On the business-level application settings page, click **Save**.
The name and description are shown in the list of applications on the Business-level applications page. Because the application is empty, its status is Unavailable.
3. Add the asset composition unit to your business-level application.
 - a. On the Business-level applications page, click the application name in the list of applications.
 - b. On the business-level application settings page, click **Add > Add Asset**.
 - c. On the Add composition unit page, select an asset composition unit from the list of available units, and then click **Continue**.
For example, select the compressed file asset and then click **Continue**.
 - d. On the Set options page, click **Next**.
 - e. On the Map composition unit to a target page, change the target server as needed, and then click **Next**.
 - f. On the Summary page, click **Finish**. Several messages are displayed. A message having the format `Completed res=[WebSphere:cuname=unit_name]` indicates that the addition is successful.
 - g. If the addition is successful, click **Manage application**.
 - h. On the business-level application settings page, click **Save**.
The asset name and type displays in the list of deployed assets.
4. Start the business-level application.
 - a. Click **Applications > Application Types > Business-level applications**.
 - b. On the Business-level applications page, select the check box beside your application.
 - c. Click **Start**.
When the business-level application is running, a green arrow displays for **Status**. If the business-level application does not start, ensure that the deployment target to which the application maps is running and try starting the application again.

What to do next

You can add other assets to your business-level application.

SCA application package deployment

The product supports deployment of many types of Service Component Architecture (SCA) artifacts as composition units of business-level applications. Typical artifacts include Java archive (JAR) files, compressed .zip files, and web application archive (WAR) files.

Details about deployment of SCA artifacts follow.

- Deployment of JAR or compressed files
- Deployment of WAR files
- Notes and limitations

Deployment of JAR or compressed files

- The product supports one composite file for each package. The product determines which composite file to support using the following process:
 1. The SCA deployment extension looks for the META-INF/sca-contribution.xml file, gets the name of each deployable composite defined in the file, and uses QName values to find the actual composite file names under any directory for that composite. If more than one composite is found in the sca-contribution.xml file, you can select the composite to deploy.
 2. If there is no META-INF/sca-contribution.xml file defined, the SCA deployment extension looks for a composite file in the META-INF/sca-deployables directory.
- The product validates SCA composites for inconsistencies with SCA specifications. One specification requirement is that the names of top-level components must be unique. Thus, the product validates top-level component name uniqueness.

Tip: Top-level components are also called domain components, with the top-level or domain typically the cell on multiple-server installations and the server scope on single-server installations.

The product validates all composite files in a JAR or compressed file, regardless of the file location in the archive or whether the sca-contribution.xml file references the composite file. The product does not validate all services and references.

The product writes warning and error messages resulting from the validation tests to the SystemOut.log file. Read the log file to learn about inconsistencies with specifications in your SCA composites.

Note: This topic references one or more of the application server log files. Beginning in WebSphere Application Server Version 8.0 you can configure the server to use the High Performance Extensible Logging (HPEL) log and trace infrastructure instead of using SystemOut.log , SystemErr.log, trace.log, and activity.log files or native z/OS logging facilities. If you are using HPEL, you can access all of your log and trace information using the LogViewer command-line tool from your server profile bin directory. See the information about using HPEL to troubleshoot applications for more information on using HPEL.

- The product uses the following process for QName resolution:
 - The product uses the QName to resolve composite files included in the top-level composite that use the element. For example, the <include name="mynamespace:MyService"/> statement looks for a composite file whose composite name is MyService and whose targetNamespace is mynamespace. The following rules apply:
 - **name:** Use the outer composite.
 - **namespace declarations:** Merged into the outer composite.
 - **targetNamespace:** Use the outer composite (must be the same).
 - **local:** Use the composite (preferably the same but not required).
 - **requires(intents) and policySets:** Must be compatible, and aggregated into the outer composite.
 - Deployable composite files must have name and targetNamespace values. The name and targetNamespace values are combined to form the QName of a composite file.
 - When a composite is used as a component implementation in the top-level composite, the composite is also resolved using the QName. For example, the <implementation.composite name="mynamespace:MyComposite"/> statement causes the product administration to look for a composite file whose composite name is MyComposite and whose targetNamespace is mynamespace.
- A JAR file can contain other JAR files at the top level. The contained JAR files are available on the classpath. However, any archives inside those JAR files are not available. The product supports one level of embedded JAR files.

Deployment of WAR files

- A composite file in a WAR file must be named default.composite. A composite file that is not in a WAR file can have any name.

- The default `composite` file must be inside a WAR file in the `META-INF/sca-deployables` directory.
- The `META-INF/sca-deployables` directory must contain no more than one composite file. If there is more than one composite file in the `META-INF/sca-deployables` directory, then the product returns a validation error and stops the WAR file deployment.
However, you can place other composite files in directories other than `META-INF/sca-deployables`, and reference those composite files in the top-level composite under the `META-INF/sca-deployables` directory.
- The product does not support having a `sca-contribution.xml` file inside the WAR file under the `META-INF` directory. If the product finds a `sca-contribution.xml` file, then the product returns a validation error and stops the WAR file deployment.

Notes and limitations

- The product does not provide administration console pages for configuring contributions.
- If you import a package or namespace from a different contribution, or JAR (`contribution.xml`), you might need to import that contribution as an asset before importing your own asset.
For example, suppose your Contribution A imports a JAR file from Contribution B. You might need to import Contribution B and then Contribution A as assets. Contribution A depends on Contribution B so you must import Contribution B before importing Contribution A.
- You cannot use a local interface across a class loader boundary. Use a remotable interface to cross a class loader boundary.

Creating SCA business-level applications

You can create an empty business-level application and then add Service Component Architecture (SCA) assets, shared libraries, business-level applications, and other artifacts as composition units to the empty business-level application.

Before you begin

Configure the target application server. You must deploy SCA composite assets of a business-level application to a Version 8.0 server or cluster (target) or to a Version 7.0 target that is enabled for the Feature Pack for SCA.

If your SCA composite or application uses Feature Pack for SCA Version 1.0.1 functionality, you must deploy the SCA asset or application to a Version 8.0 target or to a feature pack Version 1.0.1.0 target. Version 1.0.1 functionality includes:

- Java Message Service (JMS) bindings
- Atom bindings
- HTTP bindings with a wire format of JSON-RPC
- Java Platform, Enterprise Edition (Java EE) integration modules (`implementation.jee`, `implementation.web`, or `implementation.ejb` components)
- SCA Spring component implementations
- OSGi applications as SCA component implementations
- Service Data Objects (SDO) composites

Optionally, determine what assets or other files that you want to add to your business-level application and whether your application files can run on your deployment targets.

About this task

You can create business-level applications using the administrative console, the `wsadmin` tool, or programming.

You create SCA business-level applications the same way as for non-SCA business-level applications. However, when you use an SCA asset in a business-level application, function that applies only to applications that use SCA composites becomes available. For example, you can access administrative console pages that apply only to applications that use SCA composites.

Procedure

1. Select a way to create your business-level application.

Table 60. Ways to create SCA business-level applications. You can create a business-level application using the administrative console, wsadmin scripts, or programming.

Option	Method
Administrative console business-level application creation wizard	Click Applications > New Application > New Business Level Application and follow instructions in the wizard.
See "Creating SCA business-level applications with the console" on page 425.	For example use of the console to create a business-level application that has an SCA asset, see "Example: Creating an SCA business-level application with the console" on page 455.

2. Create your business-level application using the administrative console, wsadmin, or programming.
3. Save the changes to your administrative configuration.
When saving the configuration, synchronize the configuration with the nodes where the application is expected to run.

Results

The name of the application is shown in the list on the Business-level applications page.

What to do next

After you create a business-level application, you can do the following to add composition units to it:

1. Import any SCA or other assets needed by your business-level application.
2. Add assets, shared libraries, or other business-level applications as composition units.
When you add an asset, you must specify a target that supports SCA composites. Specify only a single server or cluster as the target. Do not map an SCA composition unit to multiple servers or clusters.
If the asset or application uses Feature Pack for SCA Version 1.0.1 functionality, specify a Version 8.0 target or a feature pack Version 1.0.1.0 target.
For applications that use `implementation.osgiapp`, add the enterprise bundle archive (EBA) asset as a composition unit to the business-level application before adding the SCA asset as a composition unit.
3. Save the changes to your administrative configuration.
4. Start the business-level application.

If the application does not run as desired, edit the application configuration, then save and run it again.

If the business-level application does not start, ensure that the deployment target to which the application maps is running and try starting the application again. If SCA composite assets do not start, ensure that each asset is mapped to a deployment target that supports SCA composites.

If an asset composition unit uses an Enterprise JavaBeans (EJB) binding and does not start because it has a non-WebSphere target of "null", delete the asset composition unit and add it again to the business-level application. Specify a target that supports SCA composites when you add the asset to the business-level application. You cannot change the target after deployment.

If the META-INF/sca-deployables directory has multiple SCA composite files and the application does not start because the product cannot obtain the CompUnitInfoLoader value, place only the file that contains the composite in the META-INF/sca-deployables directory. You can place the other composite files anywhere else within the archive.

If the SCA application uses security, the target must be in the global security domain.

In multiple-node environments, synchronize the nodes after you save changes to the target before starting the business-level application.

For applications that use `implementation.osgiapp` in multiple-node environments, target the EBA composition unit to the same server or cluster as the SCA composition unit.

Creating SCA business-level applications with the console

You can create an empty business-level application and then add Service Component Architecture (SCA) assets, shared libraries, or business-level applications as composition units to the empty business-level application.

Before you begin

Configure the target application server. You must deploy SCA composite assets of a business-level application to a Version 8.0 server or cluster (target) or to a Version 7.0 target that is enabled for the Feature Pack for SCA.

If your SCA composite or application uses Feature Pack for SCA Version 1.0.1 functionality, you must deploy the SCA asset or application to a Version 8.0 target or to a feature pack Version 1.0.1.0 target. Version 1.0.1 functionality includes:

- Java Message Service (JMS) bindings
- Atom bindings
- HTTP bindings with a wire format of JSON-RPC
- Java Platform, Enterprise Edition (Java EE) integration modules (`implementation.jee`, `implementation.web`, or `implementation.ejb` components)
- SCA Spring component implementations
- OSGi applications as SCA component implementations
- Service Data Objects (SDO) composites

Also, determine an application name. Optionally, determine which assets, shared libraries, or business-level applications that the new business-level application needs.

About this task

You can create a business-level application that has SCA assets using the administrative console. Alternatively, you can use the `wsadmin` scripting tool or programming.

You can add an asset or shared library composition unit to multiple business-level applications. However, each composition unit for the same asset must have a unique composition unit name. You can add a business-level application composition unit to more than one business-level application.

You must target an SCA composition unit to a single server or cluster, and not to multiple servers or clusters.

Procedure

1. Create an empty business-level application.
 - a. Click **Applications > New Application > New Business-level Application**.

- b. On the New business-level application page, specify a unique name for the application and a description, and then click **Apply**.
- c. On the business-level application settings page, click **Save**.

The name and description are shown in the list of applications on the Business-level applications page. Because the application is empty, its status is Unknown.

2. Add one SCA asset to your business-level application. The product adds the asset as a composition unit of your business-level application.
 - a. Import the SCA asset.
 - b. Go to the business-level application settings page.
Click **Applications > Application Types > Business-level applications > *application_name***.
 - c. On the business-level application settings page, specify the type of composition unit to add.
Although you can add an asset, shared library, or business-level application to your business-level application, the logic is in your SCA asset. Add the SCA asset as a composition unit.
Under **Deployed assets**, click **Add > Add Asset**.
 - d. On the Add page, select one unit from the list of available units, and then click **Continue**.
On the Add page, you might be able to select multiple deployable SCA composites. However, you can deploy only one deployable SCA composite at a time. Select only one unit and click **Continue**. If you select multiple units, the product deploys only one of those units.
For applications that use `implementation.osgiapp`, add the enterprise bundle archive (EBA) asset as a composition unit to the business-level application before adding the SCA asset as a composition unit.
 - e. On the Set options page, change the composition unit settings as needed, and then click **Next**.
This page is not shown if you have multiple deployable unit assets.
 - f. On the Map composition unit to a target page, specify one target server that supports SCA composites, and then click **Next**.
The target server can be an existing cluster. To map the composition unit to a cluster, select the existing cluster from the **Available** list, click **Add**, and then click **Next**. The cluster name is shown in the **Current targets** list as `WebSphere:cluster=cluster_name`.
You must specify only a single server or cluster as the target, and not map an SCA composition unit to multiple servers or clusters.
If you are adding an SCA asset that uses security, specify a target server that is in the global security domain.
For applications that use `implementation.osgiapp` in multiple-node environments, target the EBA composition unit to the same server or cluster as the SCA composition unit.
This page is not shown when you add a business-level application.
 - g. On the Relationship options page, click **Next** to accept the default values.
The relationships in SCA applications are set at the asset level. Either the asset must be defined as an SCA contribution or, in the asset view, a relationship must be set to another asset. When a relationship is set to another asset manually at the asset level, the relationship only exposes all the packages within the asset to the other depended asset. The namespaces are not exposed.
This page is shown only for SCA assets that have multiple deployable or composition units.
 - h. On the Set Java EE composition unit relationship page, associate SCA components with Java EE applications and then click **Next**.
Java EE applications are also known as enterprise applications or enterprise archive (EAR) files. An SCA composite definition can specify an EAR file to use on the `archive` attribute of an `implementation.jee` tag. Use this page to associate SCA components in this business-level application to the EAR files named in the composite definition. If your SCA application does not use EAR files, take the default values and click **Next**.

- i. On the Map security roles to users or groups page, specify security roles for users or groups as needed, and then click **Next**.
This page is only shown for SCA assets that use security.
 - j. On the Map RunAs roles to users page, map a user identity and password to RunAs roles as needed, and then click **Next**.
This page is only shown for SCA assets that use security.
 - k. On the Map virtual host page, specify a virtual host that hosts web services for each SCA composite, and then click **Next**. By default, composites map to `default_host`.
This page is only shown for SCA assets that contain a web service binding.
 - l. On the Attach policy set page, attach a policy set and assign policy set bindings as needed, and then click **Next**.
This page is only shown for SCA assets that use web services.
 - m. On the Summary page, click **Finish**. Several messages are displayed, indicating whether the product adds the unit to the business-level application successfully. A message having the format `Completed res=[WebSphere:cuname=unit_name]` indicates that the addition is successful. Click **Manage application**.
If the product adds the unit successfully, the name of the unit is shown in a list of deployed assets on the business-level application settings page.
If the unit addition is not successful, read the messages and add the unit again. Correct the problems noted in the messages.
 - n. On the Adding composition unit to the business-level application page, click **Save**.
3. Optional: Add one or more assets, shared libraries, or business-level applications to your business-level application.
Repeat Step 2 to add another asset or add a shared library or business-level application.

Results

A business-level application that contains the specified composition units.

What to do next

After you create the application, save the changes to your configuration and start the application as needed.

If a composite asset is deployed to a target that does not support SCA composites, the SCA composite does not start. You must deploy an SCA asset to a target that supports SCA composites.

Map virtual host settings for SCA composites

Use this page to map Service Component Architecture (SCA) composites that use a web service binding to a virtual host. You must map the composites to the virtual host that hosts the web services.

This administrative console page displays in the business-level application creation and update wizards. To view the Map virtual host page, the asset that you add to a business-level application must contain a web service binding. To view this page, do the following:

1. Import an asset that contains a web service binding.
2. Create a business-level application to which to add the asset.
3. Click **Applications > Application Types > Business-level applications > *application_name* > Add > Add Asset**.
4. On the Add composition unit page, select the asset that contains a web service binding, and click **Continue**.
5. On the Set options page, change the settings as needed and click **Next**.
6. On the Map composition unit to a target page, specify target servers as needed and click **Next**.

7. On the Define relationship with existing composition units page, change the settings as needed and click **Next**.
8. Continue changing settings as needed and click **Next** on any other pages until the Map virtual host page is displayed in the wizard.

Composite Name:

Specifies the name of the composite that uses a web service binding in the SCA artifact.

Virtual Host:

Specifies a virtual host to associate with the composite.

Select the virtual host that hosts the web services for the composite. By default, the product associates a component with the `default_host` virtual host.

Set Java EE composition unit relationships for SCA composites

Use this page to associate Service Component Architecture (SCA) components in an SCA composite with Java Platform, Enterprise Edition (Java EE) applications, otherwise known as *enterprise applications* or *enterprise archive (EAR) files*.

An SCA composite definition can define Java EE applications as component implementations. You can define an EAR asset on an `archive` attribute of the `implementation.jee` tag for the component and use the application deployed from the asset as its implementation. On this page, associate SCA components in a business-level application with the EAR files named in the composite definition.

This administrative console page displays in the business-level application creation and update wizards. To view the Set Java EE composition unit relationship page, the asset that you add to a business-level application must contain an SCA composite. To view this page, complete the following actions:

1. Import an asset that contains an SCA composite.
2. Create a business-level application to which to add the asset.
3. Click **Applications > Application Types > Business-level applications > *application_name* > Add > Add Asset**.
4. On the Add composition unit page, select the SCA composite asset and click **Continue**.
5. On the Set options page, change the settings as needed and click **Next**.
6. On the Map composition unit to a target page, specify target servers as needed and click **Next**.
7. On the Define relationship with existing composition units page, change the settings as needed and click **Next**.
8. Continue changing settings as needed and click **Next** on any other pages until the Set Java EE composition unit relationship page is displayed in the wizard.

Component Name:

Specifies the name of an SCA component in the SCA composite that you are deploying.

EAR Asset Name:

Specifies the name of the enterprise application, or EAR file, that the SCA component uses.

An EAR asset is an EAR file that has been imported as an asset. The EAR asset name must match the `archive` attribute on the `implementation.jee` tag in the SCA composite definition.

Associated Java EE Composition Unit:

Specifies the composition unit name of the EAR asset. Select the Java EE composition unit that the SCA component uses. To associate an SCA component with an EAR file, the EAR file must be a composition unit of your SCA business-level application.

If the Java EE composition unit that you want to associate with an SCA component is not in the drop-down list, import the EAR file that is named by the archive attribute on the `implementation.jee` tag of the SCA composite definition as an asset. Then, add the EAR asset as a composition unit of this business-level application.

You can use the **Import an asset** and **Add an asset** links on this page to add EAR files as assets and make them composition units of your SCA business-level application.

Attach policy set settings

Use this page to attach a policy set and assign policy set bindings for the composite defined in a Service Component Architecture (SCA) application.

This administrative console page displays in the Create new business-level application wizard. To have the Attach policy set page in the wizard, the SCA component in the asset that you add to a business-level application must use a web service binding, `binding.ws`, and the composite file or annotation must specify the intents or policy sets. To view this page, do the following:

1. Import an asset that uses a web service binding and a composite file or annotation that specifies the intents or policy sets.
2. Create a business-level application to which to add the asset.
3. Click **Applications > Application Types > Business-level applications > *application_name* > Add > Add Asset**.
4. On the Add composition unit page, select the asset that uses a web service binding, and click **Continue**.
5. On the Set options page, change the settings as needed and click **Next**.
6. On the Map composition unit to a target page, specify target servers as needed and click **Next**.
7. On the Define relationship with existing composition units page, change the settings as needed and click **Next**.
8. Continue changing settings as needed and click **Next** on any other pages until the Attach policy set page is displayed in the wizard.

To attach or detach a policy set or to assign a policy set binding, do the following:

1. Select a composite, component, service, reference, or binding from **Name**. The **Name** list is nested, indicating parent-child relationships. When you select a parent, the children are automatically selected.
2. Click the desired button.

Table 61. Button descriptions. Use the buttons to attach or detach policy sets and to assign policy set bindings.

Button	Resulting action
Attach	<p>Attaches a policy set to the selected composite, component, service, reference, or binding.</p> <p>When the Include default policy sets option is not enabled, the options for this button contain user-created policy sets only.</p> <p>When the Include default policy sets option is not enabled and no user-created policy sets exist, then there are no button options. You can select Include default policy sets to display the default policy set options.</p> <p>When the Include default policy sets option is enabled, the options for this button include both default policy sets and any user-created policy sets.</p> <p>To attach a policy set, select a composite, component, service, reference, or binding from Name and click Attach > <i>policy_set_option</i>.</p> <p>To close the menu list, click Attach.</p>

Table 61. Button descriptions (continued). Use the buttons to attach or detach policy sets and to assign policy set bindings.

Button	Resulting action
Detach Policy Set	Detaches a policy set from the selected composite, component, service, reference, or binding.
Assign Service Policy Set Binding	<p>Assigns a service policy set binding to the selected composite, component, service, reference, or binding. There are two default options:</p> <p>Default specifies to assign the default service policy set binding.</p> <p>Provider Sample specifies to assign a policy set binding that is provided with the product to the service.</p> <p>If you are deploying the composition unit to a server or cluster that belongs to a security domain, the list of policy set bindings consists of bindings that have been defined in the security domain to which the composition unit is being deployed.</p>
Assign Reference Policy Set Binding	<p>Assigns a reference policy set binding to the selected composite, component, service, reference, or binding. There are two default options:</p> <p>Default specifies to assign the default reference policy set binding.</p> <p>Client Sample specifies to assign a policy set binding that is provided with the product to the reference.</p> <p>If you are deploying the composition unit to a server or cluster that belongs to a security domain, the list of policy set bindings consists of bindings that have been defined in the security domain to which the composition unit is being deployed.</p>

Include default policy sets:

Specifies whether to include default policy sets. Default policy sets specify common quality of service (QoS) behavior for generic message format.

Before selecting this option, determine whether the default policy sets provide adequate QoS characteristics for your services.

By default, this option is not enabled.

Name:

Specifies a composite, component, service, reference, or binding in the artifact.

The **Name** list is nested, indicating parent-child relationships. When you select a parent, the children are automatically selected.

Intents:

Specifies the aggregate of the intents from the composite file and the annotations. SCA intents are used to describe the abstract policy requirements of a component.

The intents shown include any intents inherited from a parent.

Matched Policy Sets:

Specifies policy sets that potentially satisfy the intents.

You can include default policy sets by enabling the **Include default policy sets** check box. To exclude default policy sets, deselect the check box.

Attached Policy Set:

Specifies attached policy sets. If no value is shown, then the composite, component, service, reference, or binding is not attached to a policy set.

To attach a policy set, select a composite, component, service, reference, or binding and click an **Attach** option.

To detach a policy set, use **Detach Policy Set**. You can detach any policy set, including pre-attached policy sets.

Policy Set Binding:

Specifies service and reference policy set bindings. If no value is shown, then the composite, component, service, reference, or binding is not assigned to a policy set binding.

To assign a policy set binding, select a composite, component, service, reference, or binding and click an **Assign Service Policy Set Binding** or **Assign Reference Policy Set Binding** option.

To reset the bindings, select the **Default** option. For example, select **Assign Service Policy Set Binding > Default** or **Assign Reference Policy Set Binding > Default**.

Map security roles to users or groups collection for SCA composites

Use this page to view and manage mappings of security roles to users and groups that are used with the Service Component Architecture (SCA) composites.

To view this administrative console page, click **Applications > Application Types > Business-level applications > application_name > deployed_asset_composition_unit_name > Map security roles to users or groups**. This page is the same as the Map security roles to users or groups page in the Create new business-level application wizard. To view this page, your composition unit must support SCA security.

Different roles can have different security authorizations. Mapping users or groups to a role authorizes those users or groups to access applications defined by the role. Users, groups, and roles are defined when an application is installed or configured.

To map a role to a user or group, enable the Select check box beside the role name in the list and click a button. On the displayed page, specify one or more users or groups to map to the role.

Table 62. Button descriptions. Use the buttons to map security roles to users, groups, or special subjects.

Button	Resulting action
Map Users	Displays the Map users or groups page on which you can specify the users to have the selected security role.
Map Groups	Displays the Map users or groups page on which you can specify the groups to have the selected security role.

Table 62. Button descriptions (continued). Use the buttons to map security roles to users, groups, or special subjects.

Button	Resulting action
Map Special Subjects	<p>Maps special subjects according to the option that you select:</p> <p>None specifies to map none of the special subjects to the role.</p> <p>All Authenticated in Application's Realm specifies to map all of the authenticated users to a specified role. When you map all authenticated users to a specified role, all of the valid users in the current registry who have been authenticated can access resources that are protected by this role.</p> <p>All Authenticated in Trusted Realms specifies to map all of the authenticated users in the trusted realms to a specified role. This option gives all authenticated users who belong to the user registry access to the application's realm and all authenticated users who belong to user registries access to realms which are trusted by the current security domain.</p> <p>Everyone specifies to map everyone to a specified role. When you map everyone to a role, anyone can access the resources that are protected by this role and, essentially, there is no security.</p>

Role:

Specifies a security role.

Special Subjects:

Specifies which special subjects are mapped to the security role. This option applies only when an application uses multiple realms.

None Specifies to map none of the special subjects to the role.

All Authenticated in Application's Realm

Specifies to map all of the authenticated users to a specified role. When you map all authenticated users to a specified role, all of the valid users in the current registry who have been authenticated can access resources that are protected by this role.

All Authenticated in Trusted Realms

Specifies to map all of the authenticated users in the trusted realms to a specified role. All authenticated users who belong to the user registry that is mapped to the application's realm and all authenticated users who belong to user registries that are mapped to realms which are trusted by the current security domain are successfully authorized.

Everyone

Specifies to map everyone to a specified role. When you map everyone to a role, anyone can access the resources that are protected by this role and, essentially, there is no security.

To change the value, select the role, click **Map Special Subjects**, and select an option.

Users:

Lists the users that are mapped to the specified role within this application.

Users from the non-default realm are displayed as *user@realm*.

Groups:

Lists the groups that are mapped to this specified role within this application.

Map RunAs roles to users collection for SCA composites

Use this page to map a specified user identity and password to a RunAs role for a Service Component Architecture (SCA) composite. This page enables you to specify application-specific privileges for individual users to run specific tasks using another user identity.

To view this administrative console page, click **Applications > Application Types > Business-level applications > *application_name* > *deployed_asset_composition_unit_name* > Map RunAs roles to users**. This page is the same as the Map RunAs roles to users page in the Create new business-level application wizard.

To view this page, the components in your composition unit must contain predefined RunAs roles and support SCA security. RunAs roles are used by components that need to run as a particular role for recognition while interacting with another component.

Username:

Specifies a user name for the RunAs role user.

This user already maps to the role specified in the Mapping users and groups to roles page. You can map the user to its appropriate role by either mapping the user to that role directly or mapping a group that contains the user to that role. After you specify the user name and password for the user and select a RunAs role, click **Apply**.

Password:

Specifies the password for the RunAs user.

Role:

Specifies a security role for a user within this application.

The authorization policy is only enforced when security is enabled.

User:

Lists the user that is mapped to the specified role within this application.

Composition unit settings

Use this page to view composition unit settings and to change the configuration properties of a composition unit. The specific settings that are available for configuration can vary, depending upon the contents of the composition unit. For example, there are additional configuration settings if the asset contained in the composition unit is an SCA composite, or an OSGi application.

To view this administrative console page, click **Applications > Application Types > Business-level applications > *application_name* > *deployed_asset_name***. The deployed asset is a composition unit of the business-level application.

- “Settings that are common to all composition units” on page 417
- “Additional composition unit settings for SCA composites” on page 419
- “Additional composition unit settings for OSGi applications” on page 419

Settings that are common to all composition units

Name:

Specifies a logical name for the composition unit. You cannot change the name on this page.

Description:

Specifies a description for the composition unit.

Backing ID:

Specifies a unique identifier for a composition unit that is registered in the application management domain.

The identifier has the format `WebSphere:unit_typename=unit_name`. For example, for the `MyApp.jar` asset, the backing identifier might be `WebSphere:assetname=MyApp.jar`.

You cannot change the identifier on this page.

Data type	String
Units	Configuration unit identifier

Starting weight:

Specifies the order in which composition units are started when the server starts. The starting weight is like the startup order. The composition unit with the lowest starting weight is started first.

The value that you set for **Starting weight** determines the importance or weight of a composition unit within the business level application. For example, for the most important composition unit within a business-level application, specify 1 for **Starting weight**. For the next most important composition unit within the business-level application, specify 2 for **Starting weight**, and so on.

Note: Assign composition units upon which other composition units depend a lower starting weight than the dependent composition units. If a composition unit is not started and running before its dependent composition units, `java.lang.ClassNotFoundException` errors might result when you attempt to start the application or its modules.

Data type	Integer
Default	1
Range	0 to 2147483647

Start on distribution:

Specifies whether to start the composition unit when the product distributes the composition unit to other locations.

The default is not to start the composition unit.

This setting applies to asset or shared library composition units. This setting does not apply when the composition unit is a business-level application.

Data type	Boolean
Default	false

Recycle behavior on update:

Specifies whether the product restarts the composition unit after the composition unit is updated.

The default is to restart the composition unit after partial updating of the composition unit.

This setting applies to asset or shared library composition units. This setting does not apply when the composition unit is a business-level application.

Table 63. Option descriptions. Specifies whether to restart an asset or shared library composition unit.

Option	Description
ALL	Restarts the composition unit after the entire composition unit is updated
DEFAULT	Restarts the composition unit after the part of the composition unit is updated
NONE	Does not restart the composition unit after the composition unit is updated

Target mapping:

Specifies the current targets for the composition unit.

To change the deployment targets, click **Modify targets** then select a different set of deployment targets from the list of available clusters and servers.

For SCA, you must specify only a single server or cluster as the target. Do not map an SCA composition unit to multiple servers or clusters.

Note: When you change the deployment target of composition units in a business-level application, the startup order changes to the same order in which you remap composition unit targets, even if the starting weight for all composition units is set to 1. To avoid `java.lang.ClassNotFoundException` errors when attempting to start the remapped composition units, remap targets for composition units in the same order as that used to add the composition units or, after remapping, check starting weights to ensure that composition units upon which other composition units depend are started first.

Additional composition unit settings for SCA composites

SCA composite components:

Specifies the component names and component implementations of SCA composites in the application.

Table 64. Column descriptions. Provides the name of each component and the name of the class or code implementing the component.

Column	Description
Component Name	Specifies the name of a component associated with the SCA composite.
Component Implementation	Specifies the name of the class or code implementing the component.

None indicates that the SCA composite does not have defined components.

SCA composite properties:

Specifies the names and values of SCA composite properties in the application.

Table 65. Column descriptions. Provides the name and value of SCA composite properties.

Column	Description
Property Name	Specifies the name of an SCA composite property.
Property Value	Specifies the value of the property.

None indicates that the SCA composite does not have defined name-value properties.

SCA composite wires:

Specifies the sources and targets of wires in the SCA composite.

Table 66. Column descriptions. Provides the source and target of wires.

Column	Description
Wire Source	Specifies the source of a wire in the SCA composite.
Wire Target	Specifies the target of the wire.

None indicates that the SCA composite does not have defined wires.

Additional composition unit settings for OSGi applications

OSGi application deployment status:

The deployment status shows whether updates are available for the EBA asset that is contained in the composition unit. If a new version of an EBA asset is available, and all bundle downloads for the asset are complete, you can update the EBA composition unit so that the business-level application uses the latest configuration. You do not have to update the composition unit every time you update the asset.

There are four distinct deployment statuses for an EBA composition unit:

Using latest OSGi application deployment.

The composition unit is running the latest configuration of the backing asset and any CBA extensions.

New OSGi application deployment not yet available because it requires bundles that are still downloading.

The backing asset is currently undergoing a bundle version update, or bundles are downloading for a CBA extension.

New OSGi application deployment available.

The backing asset is available at a newer configuration than the configuration that is currently running in this composition unit, or a CBA extension has been added or replaced.

New OSGi application deployment cannot be applied because bundle downloads have failed.

The last bundle version update for the backing asset or CBA extension did not succeed, and therefore the newer configuration is not yet available.

If the status is “New OSGi application deployment available”, the **Update to latest deployment ...** button is available. Click this button to bring the EBA composition unit up-to-date and run the updated business-level application. If any of the updates need configuration changes, a wizard prompts you to update the configuration information.

When you save the changes to the EBA composition unit, the associated business-level application is updated to use the new configuration. If the business-level application is running, the bundle and configuration updates are applied immediately. If possible (that is, depending on the nature of the updates) the system applies the updates without restarting the application. Updates that pull in new use bundles at run time prompt a full restart of the application. Updates that pull in new provision bundles might also prompt a full application restart.

JMS binding settings for SCA composites

Use this page to view property settings for a Java Message Service (JMS) binding that connects Service Component Architecture (SCA) composite references or services. Also use this page to edit resources of the JMS binding. The settings shown on this page define a `binding.jms` element.

To view this administrative console page, your composition unit must support a JMS binding reference or service. In the administrative console, do the following:

1. Click **Applications > Application Types > Business-level applications > *application_name* > *deployed_asset_composition_unit_name***.
2. From the composition unit settings page for SCA composites, select to view references or services:
 - For an SCA component reference, click **SCA Composite Components > *SCA_component_reference_name* > SCA Component References > *reference_name* > Bindings > JMS binding**.
 - For an SCA component service, click **SCA Composite Components > *SCA_component_service_name* > SCA Component Services > *service_name* > Bindings > JMS binding**.

You can use a JMS binding to identify existing JMS resources using Java Naming and Directory Interface (JNDI) names.

- General properties (read-only)
- Resources
- Response resources
- Request header properties for a reference (read-only)
- “Response header properties” on page 440 for a service (read-only)

General properties

Specifies property settings that apply to all `binding.jms` elements. The settings pertain to both services and references. You cannot edit the settings.

JMS binding URI

Specifies a uniform resource identifier (URI) that identifies properties such as the destination, connection factory and activation specification to be used to send or receive the JMS message.

The URI has the following format:

```
jms: jms_destination?  
connectionFactoryName=connection_factory_name &  
destinationType={queue|topic} &  
deliveryMode=delivery_mode &  
timeToLive=time_to_live &  
priority=priority &  
user_property=user_property_value & ...
```

Correlation schema

Specifies the correlation scheme used when sending reply or callback messages.

Valid values are `requestmsgidtocorrelid` (the default), `requestcorrelidtocorrelid`, and `none`.

Initial context factory

Specifies the name of the initial context factory used to obtain a JNDI initial context.

This setting is optional. If no factory is specified, the WebSphere Application Server initial context factory is used.

JNDI URL

Specifies the uniform resource locator (URL) for the JNDI provider.

Request connection

Specifies a `binding.jms` element that is present in a composite definition file.

A request is a message that is sent to an SCA service or sent by an SCA reference.

Response connection

Specifies a `binding.jms` element that is present in a composite definition file.

A response is a message received by a reference (that is, a reply from an invoked service) or a message sent by a service in response to a previous request message. A response in SCA is always a reply to a previous request.

Request wire format

Specifies the component definition element name of a wire that connects SCA composites.

For example, `wireFormat.jmsObject`.

A request is a message that is sent to an SCA service or sent by an SCA reference.

Response wire format

Specifies the component definition element name of a wire that connects SCA composites.

For example, `wireFormat.jmsObject`.

A response is a message received by a reference (that is, a reply from an invoked service) or a message sent by a service in response to a previous request message. A response in SCA is always a reply to a previous request.

Resources

Specifies resources for the JMS binding. You can both read and edit resource settings for **Destination JNDI name**, **Activation specification JNDI name**, and **Connection factory JNDI name**. All other resource settings are read-only.

Destination type

Specifies the type of the request destination. Permitted values are `queue` (the default value) and `topic`. When `topic` is specified, then all the operations in the interface that correspond to the binding must be one-way.

This setting is for both services and references and is read-only.

`@type` is the destination type.

Destination JNDI name

Specifies an optional parameter that gives the JNDI name of a destination to which the binding is connected. For example, `jms/InvokeService_Callback`.

You can set a destination name for both services and references.

`@name` is the JNDI name of the destination.

Destination create

Specifies whether to create a destination for the binding.

Valid values are `ifnotexist` (the default), `always`, and `never`. When the value is `ifnotexist`, the product dynamically creates destination resources necessary for the SCA composite, if those resources do not exist and relate to the default messaging provider. The product creates the resources when adding the SCA composite to a business-level application.

The product validates a composite definition when adding an SCA asset to a business-level application. If the validation results in an error, the product does not add the asset to the application. If the value is `always` and the destination exists, deployment stops with an error.

Activation specification JNDI name

Specifies the activation specification that the binding uses to connect to a JMS destination to process request messages. The value must be a JNDI name.

You can set an activation specification for services only.

The attributes of this element follow those defined for the destination element.

Activation specification create

Specifies whether to create an activation specification for the binding. You can create an activation specification for services only.

Valid values are `ifnotexist` (the default), `always`, and `never`. When the value is `ifnotexist`, the product dynamically creates activation specification resources necessary for the SCA composite, if those resources do not exist and relate to the default messaging provider. If the value is `always` and the activation specification exists when the SCA composite is added to a business-level application, deployment stops with an error.

Connection factory JNDI name

Specifies the connection factory that the binding uses to process request messages. The value must be a JNDI name.

You can set a connection factory JNDI name for references only.

The attributes of this property follow those defined for the destination element. This property is mutually exclusive with the `activationSpec` property.

Connection factory create

Specifies whether to create a connection factory for the binding. You can create a connection factory for references only.

Valid values are `ifnotexist` (the default), `always`, and `never`. When the value is `ifnotexist`, the product dynamically creates connection factory resources necessary for the SCA composite, if those resources do not exist and relate to the default messaging provider. If the value is `always` and the connection factory exists when the SCA composite is added to a business-level application, deployment stops with an error.

Response resources

Specifies the resources used for handling response messages, receiving responses for a reference, and for sending responses from a service. A response element defines the destination and either the connection factory or activation specification elements for handling response messages.

Response resources pertain to both services and references. You can both read and edit resource settings for **Response destination JNDI name** and **Response connection factory JNDI name**. All other resource settings are read-only.

Response destination type

Specifies the type of the response destination. Permitted values are `queue` (the default value) and `topic`. When `topic` is specified, then all the operations in the interface that corresponds to the binding must be one-way.

This setting is for both services and references and is read-only.

Response destination JNDI name

Specifies the destination that is to be used to process responses by this binding. Attributes are the same as for the parent destination element. For example, `.jms/InvokeService_Response`.

You can set a destination name for both services and references.

Response destination create

Specifies whether to create a response destination for the binding.

Valid values are `ifnotexist` (the default), `always`, and `never`. When the value is `ifnotexist`, the product dynamically creates response destination resources necessary for the SCA composite, if those resources do not exist and relate to the default messaging provider. If the value is `always` and the response destination exists when the SCA composite is added to a business-level application, deployment stops with an error.

Response connection factory JNDI name

Specifies the connection factory that the binding uses to process response messages. The value must be a JNDI name.

You can set a response connection factory for both services and references.

The attributes of this element follow those defined for the destination element. This element is mutually exclusive with the activationSpec element.

Response connection factory create

Specifies whether to create a response connection factory that the binding can use to process response messages.

Valid values are `ifnotexist` (the default), `always`, and `never`. When the value is `ifnotexist`, the product dynamically creates response connection factory resources necessary for the SCA composite, if those resources do not exist and relate to the default messaging provider. If the value is `always` and the response connection factory exists when the SCA composite is added to a business-level application, deployment stops with an error.

Request header properties

Specifies JMS header properties that apply to requests from a reference. You cannot edit the header properties.

If a JMS header property is specified, the property must not appear in the URI.

JMS type

Specifies a JMS type to use in the JMS header property using `@JMSType`.

JMS correlation ID

Specifies a JMS correlation identification to use in the JMS header property using `@JMSCorrelationID`.

JMS delivery mode

Specifies a JMS delivery mode to use in the JMS header property using `@JMSDeliveryMode`.

JMS time to live

Specifies a JMS time to live to use in the JMS header property using `@JMSTimeToLive`.

JMS priority

Specifies a JMS priority to use in the JMS header property using `@JMSPriority`.

Header property

Specifies a value to use for the specified JMS user property.

Response header properties

Specifies JMS header properties that apply to responses from a service for outbound messages. You cannot edit the header properties.

If a JMS header property is specified, the property must not appear in the URI.

JMS type

Specifies a JMS type to use in the JMS header property using `@JMSType`.

JMS correlation ID

Specifies a JMS correlation identification to use in the JMS header property using `@JMSCorrelationID`.

JMS delivery mode

Specifies a JMS delivery mode to use in the JMS header property using `@JMSDeliveryMode`.

JMS time to live

Specifies a JMS time to live to use in the JMS header property using `@JMSTimeToLive`.

JMS priority

Specifies a JMS priority to use in the JMS header property using `@JMSPriority`.

Header property

Specifies a value to use for the specified JMS user property.

Provide HTTP endpoint URL information settings for SCA composites

Use this page to specify endpoint Universal Resource Locator (URL) prefix information for Service Component Architecture (SCA) composites accessed by web service bindings. The information is used to form complete endpoint addresses.

To view this administrative console page, click **Applications > Application Types > Business-level applications > *application_name* > *deployed_asset_composition_unit_name* > Provide HTTP endpoint URL information**.

Default SCA URL prefixes

Shows the predefined default endpoint URL prefixes for SCA composites that are accessed by Hypertext Transfer Protocol (HTTP) or Hypertext Transfer Protocol Secure (HTTPS) for service endpoints.

The field shows both unsecure and secure custom endpoint URL values separated by a space. For example:

```
http://theHost:9081 https://theHost:9044
```

For each endpoint URL prefix, the format is *protocol://host_name:port_number*. The protocol is either `http` or `https`. In this example, *host_name* is `theHost` and *port_number* is the port used in the endpoint URL.

To use the default endpoint URL prefixes, deselect the **Override default SCA URL prefixes** check box.

Override default SCA URL prefixes

Specifies whether to use a custom endpoint URL prefix when the service has a proxied front end. The endpoint URL prefixes are those of the proxy server. You must specify proxied endpoints when deploying services that use the web service binding in a clustered configuration.

To specify a custom endpoint URL prefix, do the following:

1. Select the **Override default SCA URL prefixes** check box.
2. For **HTTP host name**, specify the host name of the unsecure custom endpoint. For example: `myHost`
3. For **HTTP port**, specify the port of the unsecure custom endpoint. For example: `9081`
4. For **HTTPS host name**, specify the host name of the secure custom endpoint. For example: `myHost`
5. For **HTTPS port**, specify the port of the secure custom endpoint. For example: `9044`
6. Click **OK**.

SCA composite component settings

Use this page to view and edit the attributes associated with a Service Component Architecture (SCA) component.

To view this administrative console page, click **Applications > Application Types > Business-level applications > *application_name* > *deployed_asset_name* > *SCA_composite_component_name***.

Components are configured instances of implementations. Components provide and consume services. More than one component can use and configure the same implementation, where each component configures the implementation differently. For example each component might configure a reference of the same implementation to consume a different service.

An implementation defines the aspects configurable by a component in the form of a component type. The component type is in effect a description of the contract honored by the implementation.

A reference represents a requirement that the implementation has on a service provided by another component.

Component name

Specifies the component name of the attribute.

Implementation

Specifies the name of the class or configuration file that contains the component implementation.

For `implementation.java`, the Java class is shown. For other implementations, the name of the resource identified by the implementation is shown:

- For `implementation.jee`, the archive name is shown.
- For `implementation.spring`, the application context file is shown.
- For `implementation.osgiapp`, the application symbolic name and version is shown.

Type

Specifies the type of attribute. In this case, the type is Component.

SCA component services

Specifies the names of the services.

SCA component references

Specifies the names and targets of component references. You can edit the reference target for customization.

SCA component properties

Specifies the **Property Input Source** and **Property Value** for each property.

Options for **Property Input Source** include the following:

- **XPath** indicates the source attribute of the property.
- **File** indicates the file attribute of the property.
- **Value** indicates the property element value.

SCA component reference settings

Use this page to view and edit the attributes associated with a Service Component Architecture (SCA) component reference.

To view this administrative console page, click **Applications > Application Types > Business-level applications > *application_name* > *deployed_asset_composition_unit_name* > *SCA_composite_component_name* > *reference_name***.

SCA component references within an implementation represent links to services the implementation uses that must be provided by other components in the SCA system. For a composite, you can wire references of components within the composite (component references) to references of the composite (composite references), indicating that the component references must be resolved by services outside the composite.

References use bindings to describe the access methods used to invoke the services.

Under **Additional Properties**, click **View domain** to view a list of services available in the current cell or domain. This can be helpful when updating the **Target** setting value, for example.

Reference name

Specifies the reference name of the attribute.

Type

Specifies the type of attribute. In this case, it is Reference.

Reference target URI

Specifies one or more target service uniform resource identifiers (URIs), depending on the multiplicity setting. Each target wires the reference to a component service that resolves the reference. Targets can contain a list of targets separated by a space, in the form **target1 target2**.

Bindings

Specifies the URI of the binding.

Supported bindings include the SCA default binding, enterprise bean (EJB) binding, web service binding, Java Message Service (JMS) binding, Atom binding, and HTTP binding.

SCA component service settings

Use this page to view and edit the attributes associated with a component service.

To view this administrative console page, click **Applications > Application Types > Business-level applications > application_name > deployed_asset_composition_unit_name > SCA_composite_component_name > service_name**.

Services are used to publish services provided by implementations, so that they are addressable by other components.

A service published by a component can be provided by a service of a component defined within the component, or it can be provided by a component reference. The latter case allows the republication of a service with a new address or new bindings.

Service name

Specifies the service name of the attribute.

Type

Specifies the type of attribute. In this case, *Service*.

Work manager JNDI name

Specifies the Java Naming and Directory Interface (JNDI) name of the work manager.

Bindings

Specifies the uniform resource identifier (URI) of the binding.

Supported bindings include the SCA default binding, enterprise bean (EJB) binding, web service binding, Java Message Service (JMS) binding, Atom binding, and HTTP binding.

Service provider policy sets and bindings collection for SCA composites

Use this page to attach and detach policy sets to a composition unit, a service provider, its endpoints, or operations of a Service Component Architecture (SCA) composite. You can select the default bindings, create new application-specific bindings, or use bindings that you created for an attached policy set. You can view or change whether the service provider can share its current policy configuration.

To view this administrative console page, your composition unit must use web services and support SCA. Click **Applications > Application Types > Business-level applications > application_name > deployed_asset_composition_unit_name > Service provider policy sets and bindings** .

Depending on your assigned security role when security is enabled, you might not have access to text entry fields or buttons to create or edit configuration data. Review the administrative roles documentation to learn more about the valid roles for the application server.

To attach or detach a policy set or binding, do the following:

1. Select a composition unit, service, endpoint, or operation. The **Composition unit/Service/Endpoint/Operation** list is nested, indicating parent-child relationships.
2. Click the desired button.

Table 67. Button descriptions. Use the buttons to attach or detach policy sets and to assign policy set bindings.

Button	Resulting action
Attach	<p>Attaches a policy set to the selected composition unit, service, endpoint, or operation. To attach a policy set, select a unit, service, endpoint, or operation and click Attach > policy_set_option.</p> <p>To close the menu list, click Attach.</p>
Detach Policy Set	<p>Detaches a policy set from the selected composition unit, service, endpoint, or operation.</p> <p>After the policy set is detached, if there is no policy set attached to an upper-level service resource, the Attached Policy Set column displays None and the Binding column displays Not applicable.</p> <p>If there is a policy set attached to an upper-level service resource, the Attached Policy Set column displays policy_set_name (inherited) and the binding used for the upper-level attachment is applied. The binding name is displayed followed by (inherited).</p>

Table 67. Button descriptions (continued). Use the buttons to attach or detach policy sets and to assign policy set bindings.

Button	Resulting action
Assign Binding	<p>Assigns a policy set binding to the selected composition unit, service, endpoint, or operation. The options include the following:</p> <p>Default Specifies the default binding for the selected service reference, endpoint, or operation. You can specify client and provider default bindings to be used at the cell level or global security domain level, for a particular server, or for a security domain. The default bindings are used when an application-specific binding has not been assigned to the attachment. When you attach a policy set to a service resource, the binding is initially set to the default. If you do not specifically assign a binding to the attachment point using this Assign Binding action, the default specified at the nearest scope is used.</p> <p>For any policy set attachment, the run time checks to see if the attachment includes a binding. If so, it uses that binding. If not, the run time checks in the following order and uses the first available default binding:</p> <ol style="list-style-type: none"> 1. Default general bindings for the server 2. Default general bindings for the domain in which the server resides 3. Default general bindings for the global security domain <p>New Application Specific Binding Select this option to create a new application-specific binding for the policy set attachments. The new binding you create is used for the selected resources. If you select more than one resource, ensure that all selected resources have the same policy set attached.</p> <p>Provider sample Select this option to use the Provider sample binding.</p> <p>Provider sample V2 Select this option to use the Provider sample V2 binding when you are using either the Kerberos V5 WSSecurity default or the TrustServiceKerberosDefault policy sets.</p> <p>Saml Bearer Provider sample Select this option to use the Saml Bearer Provider sample. The Saml Bearer Provider sample extends the Provider sample binding to support SAML Bearer token usage scenarios. You can use this sample with any of the SAML bearer token default policy sets.</p> <p>Saml HoK Symmetric Provider sample Select this option to use the Saml HoK Symmetric Provider sample. The Saml HoK Symmetric Provider sample extends the Provider sample binding to support SAML holder-of-key (HoK) symmetric key token usage scenarios. You can use this sample with one of the SAML HoK Symmetric key default policy sets: either SAML11 HoK Symmetric WSSecurity default or SAML20 HoK Symmetric WSSecurity default.</p>

Composition unit/Service/Endpoint/Operation

Specifies the name of the composition unit and the associated service providers, endpoints or operations.

The Composition unit/Service/Endpoint/Operation column lists the composition unit and the service providers, endpoints, or operations that the composition unit contains.

Attached Policy Set

Specifies the policy set that is attached to a composition unit, service provider, endpoint, or operation.

The Attached Policy Set column can contain the following values:

- **None.** No policy set is attached, either directly or to a higher-level service resource.
- ***Policy_set_name***. The name of the policy set that is attached directly to the service resource, for example, WS-I RSP.

- **Policy_set_name (inherited).** The name of the policy set that is not attached directly to a service resource, but that is attached to a higher-level service resource.

When the value in the column is a link, click the link to view or change settings about the attached policy set.

Binding

Specifies the binding configuration that is available for a service provider, endpoint, or operation.

The Binding column can contain the following values:

- **Not applicable.** No policy set is attached, either directly or to a higher-level service resource.
- **Binding_name** or **Default.** The binding name is displayed if a policy set is attached directly and an application-specific binding or a general binding is assigned, for example, MyBindings1. **Default** is displayed if a policy set is attached directly but the service resource uses the default bindings.
- **Binding_name (inherited)** or **Default (inherited).** A service resource inherits the bindings from an attachment to a higher-level resource.

When the value in the Binding column is a link, click the link to view or change settings about the binding.

References policy sets and bindings collection for SCA composites

Use this page to attach and detach policy sets to a composition unit, a service reference, its endpoints, or operations of a Service Component Architecture (SCA) composite. You can select the default bindings, create new application-specific bindings, or use bindings that you created for an attached policy set. You can view or change whether the service reference can share its current policy configuration.

To view this administrative console page, your composition unit must use web services and support SCA. Click **Applications > Application Types > Business-level applications > application_name > deployed_asset_composition_unit_name > References policy sets and bindings.**

Depending on your assigned security role when security is enabled, you might not have access to text entry fields or buttons to create or edit configuration data. Review the administrative roles documentation to learn more about the valid roles for the application server.

To attach or detach a policy set or binding, do the following:

1. Select a composition unit, service, endpoint, or operation. The **Composition unit/Service/Endpoint/Operation** list is nested, indicating parent-child relationships.
2. Click the desired button.

Table 68. Button descriptions. Use the buttons to attach or detach client policy sets and to assign policy set bindings.

Button	Resulting action
Attach Client Policy Set	Attaches a client policy set to the selected composition unit, service, endpoint, or operation. To attach a policy set, select a composition unit, service, endpoint, or operation and click Attach Client Policy Set > policy_set_option. To close the menu list, click Attach Client Policy Set.

Table 68. Button descriptions (continued). Use the buttons to attach or detach client policy sets and to assign policy set bindings.

Button	Resulting action
Detach Client Policy Set	<p>Detaches a client policy set from the selected composition unit, service, endpoint, or operation.</p> <p>After the policy set is detached, if there is no policy set attached to an upper-level service resource, the Attached Client Policy Set column displays None and the Binding column displays Not applicable.</p> <p>If there is a policy set attached to an upper-level service resource, the Attached Client Policy Set column displays <i>policy_set_name (inherited)</i> and the binding used for the upper-level attachment is applied. The binding name is displayed followed by (inherited).</p>
Assign Binding	<p>Assigns a policy set binding to the selected composition unit, service, endpoint, or operation. The options include the following:</p> <p>Default Specifies the default binding for the selected service, endpoint, or operation. You can specify client and provider default bindings to be used at the cell level or global security domain level, for a particular server, or for a security domain. The default bindings are used when an application-specific binding has not been assigned to the attachment. When you attach a policy set to a service resource, the binding is initially set to the default. If you do not specifically assign a binding to the attachment point using this Assign Binding action, the default specified at the nearest scope is used.</p> <p>For any policy set attachment, the run time checks to see if the attachment includes a binding. If so, it uses that binding. If not, the run time checks in the following order and uses the first available default binding:</p> <ol style="list-style-type: none"> 1. Default general bindings for the server 2. Default general bindings for the domain in which the server resides 3. Default general bindings for the global security domain <p>New Application Specific Binding Select this option to create a new application-specific binding for the policy set attachments. The new binding you create is used for the selected resources. If you select more than one resource, ensure that all selected resources have the same policy set attached.</p> <p>Client sample Select this option to use the Client sample binding.</p> <p>Client sample V2 Select this option to use the Client sample V2 binding when you are using either the Kerberos V5 WSSecurity default or the TrustServiceKerberosDefault policy sets.</p> <p>Saml Bearer Client sample Select this option to use the Saml Bearer Client sample. The Saml Bearer Client sample extends the Client sample binding to support SAML Bearer token usage scenarios. You can use this sample with any of the SAML bearer token default policy sets.</p> <p>Saml HoK Symmetric Client sample Select this option to use the Saml HoK Symmetric Client sample. The Saml HoK Symmetric Client sample extends the Client sample binding to support SAML holder-of-key (HoK) symmetric key token usage scenarios. You can use this sample with one of the SAML HoK Symmetric key default policy sets: either SAML11 HoK Symmetric WSSecurity default or SAML20 HoK Symmetric WSSecurity default.</p>

Composition unit/Service/Endpoint/Operation

Specifies the name of the composition unit and the associated service references, endpoints or operations.

The Composition unit/Service/Endpoint/Operation column lists the service composition unit and the service references, endpoints, or operations that the composition unit contains.

Attached Client Policy Set

Specifies the policy set that is attached to a composition unit, service reference, endpoint, or operation.

The Attached Client Policy Set column can contain the following values:

- **None.** No policy set is attached, either directly or to a higher-level service resource.
- ***Policy_set_name*.** The name of the policy set that is attached directly to the service resource, for example, WS-I RSP.
- ***Policy_set_name (inherited)*.** The name of the policy set that is not attached directly to a service resource, but that is attached to a higher-level service resource.

When the value in the column is a link, click the link to view or change settings about the attached policy set.

Binding

Specifies the binding configuration that is available for a service reference, endpoint, or operation.

The Binding column can contain the following values:

- **Not applicable.** No policy set is attached, either directly or to a higher-level service resource.
- ***Binding_name*** or **Default.** The binding name is displayed if a policy set is attached directly and an application-specific binding or a general binding is assigned, for example, MyBindings1. **Default** is displayed if a policy set is attached directly but the service resource uses the default bindings.
- ***Binding_name (inherited)*** or **Default (inherited).** A service resource inherits the bindings from an attachment to a higher-level resource.

When the value in the Binding column is a link, click the link to view or change settings about the binding.

SCA service provider settings

Use this page to manage policy sets for a Service Component Architecture (SCA) web service provider. You can attach and detach policy sets to a service provider, its endpoints, or operations. You can select the default bindings, create new application-specific bindings, or use bindings that you created for an attached policy set. You can view or change whether the service provider can share its current policy configuration.

To view this administrative console page, your composition unit must use web services and support SCA. Click **Services > Service providers > *service_provider_name***.

Service provider

Specifies the full QName of the service provider. The QName must be in a format that supports the Java class `javax.xml.namespace.QName`.

For the SCA sample business-level application HelloWorldAsync, the service provider name resembles the following:

```
{http://websphere.ibm.com/HelloWorldServiceComponent/HelloWorldService}HelloWorldService
```

Policy Set Attachments

Specifies the attached policy sets and assigned bindings for services, endpoints, or operations in the service provider.

To attach or detach a policy set or to assign bindings with system-specific configurations, do the following:

1. Select a service, endpoint, or operation. The **Service/Endpoint/Operation** list is nested, indicating parent-child relationships.
2. Click the desired button.

Table 69. Button descriptions. Use the buttons to attach or detach policy sets and to assign policy set bindings.

Button	Resulting action
Attach	<p>Attaches a policy set to the selected service, endpoint, or operation. To attach a policy set, select a service, endpoint, or operation and click Attach > <i>policy_set_option</i>.</p> <p>To close the menu list, click Attach.</p>
Detach Policy Set	<p>Detaches a policy set from the selected service, endpoint, or operation.</p> <p>After the policy set is detached, if there is no policy set attached to an upper-level service resource, the Attached Policy Set column displays None and the Binding column displays Not applicable.</p> <p>If there is a policy set attached to an upper-level service resource, the Attached Policy Set column displays <i>policy_set_name (inherited)</i> and the binding used for the upper-level attachment is applied. The binding name is displayed followed by (inherited).</p>

Table 69. Button descriptions (continued). Use the buttons to attach or detach policy sets and to assign policy set bindings.

Button	Resulting action
Assign Binding	<p>Assigns a policy set binding to the selected service, endpoint, or operation. The options include the following:</p> <p>Default Specifies the default binding for the selected service reference, endpoint, or operation. You can specify client and provider default bindings to be used at the cell level or global security domain level, for a particular server, or for a security domain. The default bindings are used when an application-specific binding has not been assigned to the attachment. When you attach a policy set to a service resource, the binding is initially set to the default. If you do not specifically assign a binding to the attachment point using this Assign Binding action, the default specified at the nearest scope is used.</p> <p>For any policy set attachment, the run time checks to see if the attachment includes a binding. If so, it uses that binding. If not, the run time checks in the following order and uses the first available default binding:</p> <ol style="list-style-type: none"> 1. Default general bindings for the server 2. Default general bindings for the domain in which the server resides 3. Default general bindings for the global security domain <p>New Application Specific Binding Select this option to create a new application-specific binding for the policy set attachments. The new binding you create is used for the selected resources. If you select more than one resource, ensure that all selected resources have the same policy set attached.</p> <p>Provider sample Select this option to use the Provider sample binding.</p> <p>Provider sample V2 Select this option to use the Provider sample V2 binding when you are using either the Kerberos V5 WSSecurity default or the TrustServiceKerberosDefault policy sets.</p> <p>Saml Bearer Provider sample Select this option to use the Saml Bearer Provider sample. The Saml Bearer Provider sample extends the Provider sample binding to support SAML Bearer token usage scenarios. You can use this sample with any of the SAML bearer token default policy sets.</p> <p>Saml HoK Symmetric Provider sample Select this option to use the Saml HoK Symmetric Provider sample. The Saml HoK Symmetric Provider sample extends the Provider sample binding to support SAML holder-of-key (HoK) symmetric key token usage scenarios. You can use this sample with one of the SAML HoK Symmetric key default policy sets: either SAML11 HoK Symmetric WSSecurity default or SAML20 HoK Symmetric WSSecurity default.</p>

Depending on your assigned security role when security is enabled, you might not have access to text entry fields or buttons to create or edit configuration data. Review the administrative roles documentation to learn more about the valid roles for the application server.

Service/Endpoint/Operation

Specifies the name of the service and the associated service providers, endpoints or operations.

The Service/Endpoint/Operation column lists the service and the service providers, endpoints, or operations that the service contains.

Attached Policy Set

Specifies the policy set that is attached to a service provider, endpoint, or operation.

The Attached Policy Set column can contain the following values:

- **None.** No policy set is attached, either directly or to a higher-level service resource.

- **Policy_set_name.** The name of the policy set that is attached directly to the service resource, for example, WS-I RSP.
- **Policy_set_name (inherited).** The name of the policy set that is not attached directly to a service resource, but that is attached to a higher-level service resource.

When the value in the column is a link, click the link to view or change settings about the attached policy set.

Binding

Specifies the binding configuration that is available for a service provider, endpoint, or operation.

The Binding column can contain the following values:

- **Not applicable.** No policy set is attached, either directly or to a higher-level service resource.
- **Binding_name** or **Default.** The binding name is displayed if a policy set is attached directly and an application-specific binding or a general binding is assigned, for example, MyBindings1. **Default** is displayed if a policy set is attached directly but the service resource uses the default bindings.
- **Binding_name (inherited)** or **Default (inherited).** A service resource inherits the bindings from an attachment to a higher-level resource.

When the value in the Binding column is a link, click the link to view or change settings about the binding.

About policy set bindings

In this release, there are two types of bindings: application-specific bindings and general bindings. Composition units can use both application-specific bindings and general bindings.

Application-specific bindings

You can create application-specific bindings only at a policy set attachment point. These bindings are specific to, and constrained by, the characteristics of the defined policy. Application-specific bindings can provide configuration for advanced policy requirements such as multiple signatures; however, these bindings are reusable only within an application. Also, application-specific bindings have very limited reuse across policy sets.

When you create an application-specific binding for a policy set attachment, the binding begins in a completely unconfigured state. You must add each policy, such as WS-Security or HTTP transport, that you want to override the default binding, and fully configure the bindings for each policy that you add. For WS-Security policy, some high level configuration attributes such as TokenConsumer, TokenGenerator, SigningInfo, or EncryptionInfo might be obtained from the default bindings if they are not configured in the application-specific bindings.

For service providers, you can create application-specific bindings only by selecting **Assign Binding > New Application Specific Binding**, on the Service providers policy sets and bindings collection page, for service provider resources that have an attached policy set. Similarly, for service clients, you can create application-specific bindings only by selecting **Assign Binding > New Application Specific Binding**, on the Service clients policy sets and bindings collection page, for service client resources that have an attached policy set.

General bindings

You can configure general bindings to be used across a range of policy sets and they can be reused across applications and for trust service attachments. Although general bindings are highly reusable, they cannot provide configuration for advanced policy requirements such as multiple signatures. There are two types of general bindings: general provider policy set bindings and general client policy set bindings.

You can create general provider policy set bindings by clicking **Services > Policy sets > General provider policy set bindings > New** in the general provider policy sets panel, or by clicking **Services > Policy sets > General client policy set bindings > New** in the general client policy set and bindings panel. For details about defining and managing service client or provider bindings, see the related links. General provider policy set bindings might also be used for trust service attachments.

SCA service client settings

Use this page to manage policy sets for a Service Component Architecture (SCA) web service client. You can attach and detach policy sets to a service reference, its endpoints, or operations. You can select the default bindings, create new application-specific bindings, or use bindings that you created for an attached policy set. You can view or change whether the service reference can share its current policy configuration.

To view this administrative console page, your composition unit must use web services and support SCA. Click **Services > Service clients > *service_client_name***.

Service client

Specifies the full QName of the service client. The QName must be in a format that supports the Java class `javax.xml.namespace.QName`.

For the SCA sample business-level application HelloWorldAsync, the service client name resembles the following:

```
{http://websphere.ibm.com/HelloWorldServiceComponent/HelloWorldService}HelloWorldCallbackService
```

This SCA application has the product web service namespace, `http://websphere.ibm.com/`, and the service name in its service client name.

Policy Set Attachments

Specifies the attached policy sets and assigned bindings for services, endpoints, or operations in the service client.

To attach or detach a policy set or to assign bindings with system-specific configurations, do the following:

1. Select a service, endpoint, or operation. The **Service/Endpoint/Operation** list is nested, indicating parent-child relationships.
2. Click the desired button.

Table 70. Button descriptions. Use the buttons to attach or detach client policy sets and to assign policy set bindings.

Button	Resulting action
Attach Client Policy Set	<p>Attaches a client policy set to the selected service, endpoint, or operation. To attach a policy set, select a service, endpoint, or operation and click Attach Client Policy Set > <i>policy_set_option</i>.</p> <p>To close the menu list, click Attach Client Policy Set.</p>
Detach Client Policy Set	<p>Detaches a client policy set from the selected service, endpoint, or operation.</p> <p>After the policy set is detached, if there is no policy set attached to an upper-level service resource, the Attached Client Policy Set column displays None and the Binding column displays Not applicable.</p> <p>If there is a policy set attached to an upper-level service resource, the Attached Client Policy Set column displays <i>policy_set_name</i> (inherited) and the binding used for the upper-level attachment is applied. The binding name is displayed followed by (inherited).</p>

Table 70. Button descriptions (continued). Use the buttons to attach or detach client policy sets and to assign policy set bindings.

Button	Resulting action
Assign Binding	<p>Assigns a policy set binding to the selected service, endpoint, or operation. The options include the following:</p> <p>Default Specifies the default binding for the selected service, endpoint, or operation. You can specify client and provider default bindings to be used at the cell level or global security domain level, for a particular server, or for a security domain. The default bindings are used when an application-specific binding has not been assigned to the attachment. When you attach a policy set to a service resource, the binding is initially set to the default. If you do not specifically assign a binding to the attachment point using this Assign Binding action, the default specified at the nearest scope is used.</p> <p>For any policy set attachment, the run time checks to see if the attachment includes a binding. If so, it uses that binding. If not, the run time checks in the following order and uses the first available default binding:</p> <ol style="list-style-type: none"> 1. Default general bindings for the server 2. Default general bindings for the domain in which the server resides 3. Default general bindings for the global security domain <p>New Application Specific Binding Select this option to create a new application-specific binding for the policy set attachments. The new binding you create is used for the selected resources. If you select more than one resource, ensure that all selected resources have the same policy set attached.</p> <p>Client sample Select this option to use the Client sample binding.</p> <p>Client sample V2 Select this option to use the Client sample V2 binding when you are using either the Kerberos V5 WSSecurity default or the TrustServiceKerberosDefault policy sets.</p> <p>Saml Bearer Client sample Select this option to use the Saml Bearer Client sample. The Saml Bearer Client sample extends the Client sample binding to support SAML Bearer token usage scenarios. You can use this sample with any of the SAML bearer token default policy sets.</p> <p>Saml HoK Symmetric Client sample Select this option to use the Saml HoK Symmetric Client sample. The Saml HoK Symmetric Client sample extends the Client sample binding to support SAML holder-of-key (HoK) symmetric key token usage scenarios. You can use this sample with one of the SAML HoK Symmetric key default policy sets: either SAML11 HoK Symmetric WSSecurity default or SAML20 HoK Symmetric WSSecurity default.</p>

Depending on your assigned security role when security is enabled, you might not have access to text entry fields or buttons to create or edit configuration data. Review the administrative roles documentation to learn more about the valid roles for the application server.

Service/Endpoint/Operation

Specifies the name of the service and the associated service references, endpoints or operations.

The Service/Endpoint/Operation column lists the service and the service references, endpoints, or operations that the service contains.

Attached Client Policy Set

Specifies the policy set that is attached to a service reference, endpoint, or operation.

The Attached Client Policy Set column can contain the following values:

- **None.** No policy set is attached, either directly or to a higher-level service resource.
- ***Policy_set_name*.** The name of the policy set that is attached directly to the service resource, for example, WS-I RSP.
- ***Policy_set_name (inherited)*.** The name of the policy set that is not attached directly to a service resource, but that is attached to a higher-level service resource.

When the value in the column is a link, click the link to view or change settings about the attached policy set.

Binding

Specifies the binding configuration that is available for a service reference, endpoint, or operation.

The Binding column can contain the following values:

- **Not applicable.** No policy set is attached, either directly or to a higher-level service resource.
- ***Binding_name* or *Default*.** The binding name is displayed if a policy set is attached directly and an application-specific binding or a general binding is assigned, for example, MyBindings1. **Default** is displayed if a policy set is attached directly but the service resource uses the default bindings.
- ***Binding_name (inherited)* or *Default (inherited)*.** A service resource inherits the bindings from an attachment to a higher-level resource.

When the value in the Binding column is a link, click the link to view or change settings about the binding.

About policy set bindings

In this release, there are two types of bindings: application-specific bindings and general bindings. Composition units can use both application-specific bindings and general bindings.

Application-specific bindings

You can create application-specific bindings only at a policy set attachment point. These bindings are specific to, and constrained by, the characteristics of the defined policy. Application-specific bindings can provide configuration for advanced policy requirements such as multiple signatures; however, these bindings are reusable only within an application. Also, application-specific bindings have very limited reuse across policy sets.

When you create an application-specific binding for a policy set attachment, the binding begins in a completely unconfigured state. You must add each policy, such as WS-Security or HTTP transport, that you want to override the default binding, and fully configure the bindings for each policy that you add. For WS-Security policy, some high level configuration attributes such as TokenConsumer, TokenGenerator, SigningInfo, or EncryptionInfo might be obtained from the default bindings if they are not configured in the application-specific bindings.

For service providers, you can create application-specific bindings only by selecting **Assign Binding > New Application Specific Binding**, on the Service providers policy sets and bindings collection page, for service provider resources that have an attached policy set. Similarly, for service clients, you can create application-specific bindings only by selecting **Assign Binding > New Application Specific Binding**, on the Service clients policy sets and bindings collection page, for service client resources that have an attached policy set.

General bindings

You can configure general bindings to be used across a range of policy sets and they can be reused across applications and for trust service attachments. Although general bindings are highly reusable, they

cannot provide configuration for advanced policy requirements such as multiple signatures. There are two types of general bindings: general provider policy set bindings and general client policy set bindings.

You can create general provider policy set bindings by clicking **Services > Policy sets > General provider policy set bindings > New** in the general provider policy sets panel, or by clicking **Services > Policy sets > General client policy set bindings > New** in the general client policy set and bindings panel. For details about defining and managing service client or provider bindings, see the related links. General provider policy set bindings might also be used for trust service attachments.

Example: Creating an SCA business-level application with the console

You can add many different types of artifacts to business-level applications. For example, you can add applications or modules, Java archives (JAR files), data in compressed files, and other business-level applications. This example describes how to create an empty business-level application and then add a Service Component Architecture (SCA) JAR file to the application using the administrative console.

Before you begin

In a product installation, verify that the target server is configured. As part of configuring the server, determine whether your application files can run on your deployment target. You must deploy SCA composite assets of a business-level application to a Version 8.0 server or cluster (target) or to a Version 7.0 target that is enabled for the Feature Pack for SCA.

Download the `helloworld-ws-async.jar` SCA sample file from a product download site:

1. Go to the **Samples, Version 8.0** information center.
2. On the **Downloads** tab, click **FTP** or **HTTP** in the **Service Component Architecture** section.
3. In the authentication window, click **OK**.
4. In the **SCA.zip** compressed file, go to the `SCA/installableApps` directory and download the `helloworld-ws-async.jar` file.

About this task

For this example, use the administrative console to create a business-level application named `HelloWorldAsync` that has an SCA JAR file, `helloworld-ws-async.jar`, as an asset.

Procedure

1. Create an empty business-level application named `HelloWorldAsync`.
 - a. Click **Applications > New Application > New Business Level Application**.
 - b. On the New application page, specify the name `HelloWorldAsync`, optionally add a description, and then click **Apply**.
 - c. On the page that is displayed, click the **Save** link.

The name is shown in the list of applications on the Business-level applications page. Because the application is empty, its status is `Unknown`.

2. Import the SCA JAR asset.
 - a. Click **Applications > New Application > New Asset** in the console navigation tree.
 - b. On the Upload asset page, specify the asset package to import, `helloworld-ws-async.jar`, and click **Next**.

The JAR file is in the `app_server_root/installableApps` directory.
 - c. On the Select options for importing an asset page, click **Next** to accept the default values.
 - d. On the Summary page, click **Finish**.
 - e. On the Adding asset to repository page, if messages show that the operation completed, click **Manage assets**.

- f. On the Assets page, click the **Save** link.

The file name displays in the list of assets.

3. Add the SCA JAR asset as a composition unit of the business-level application.
 - a. Click **Applications > Application Types > Business-level applications**.
 - b. On the Business-level applications page, click the HelloWorldAsync application name.
 - c. On the business-level application settings page, click **Add > Add Asset**.
 - d. On the Add page, select the helloworld-ws-async.jar asset composition unit from the list of available units, and then click **Continue**.
 - e. On the Set options page, click **Next** to accept the default values.
 - f. On the Map composition unit to a target page, specify a target server that supports SCA composites, and then click **Next**.

The target server can be an existing cluster. To map the composition unit to a cluster, select the existing cluster from the **Available** list, click **Add**, and then click **Next**. The cluster name is shown in the **Current targets** list as `WebSphere:cluster=cluster_name`.

- g. On the Define relationship with existing composition units page, click **Next** to accept the default values.
- h. On the Map virtual host page, click **Next** to accept the default values.
- i. On the Summary page, click **Finish**.

Several messages are displayed. A message having the format `Completed res=[WebSphere:cuname=helloworldws]` indicates that the addition is successful.

During deployment of the composition unit, you can view the Uniform Resource Identifier (URI) for composite level service of some bindings, along with the service name and binding type. Only the URI is editable. The product does not validate the URI.

- j. If the addition is successful, click **Manage application**.
- k. On the business-level application settings page, click **Save**.

The asset name and type displays in the list of deployed assets. If you click on the asset name, the composition unit settings page displays, with the asset name in the **SCA Composite Components** list.

4. Start the HelloWorldAsync business-level application.
 - a. Click **Applications > Application Types > Business-level applications**.
 - b. On the Business-level applications page, select the check box beside HelloWorldAsync.
 - c. Click **Start**.

When the business-level application is running, a green arrow displays for **Status**. If the business-level application does not start, ensure that the deployment target to which the application maps is running and try starting the application again.

What to do next

Optionally examine, and possibly use in applications, other SCA sample files in the `SCA/installableApps` directory of the downloadable SCA samples.

If the business-level application does not start, ensure that the deployment target to which the application maps is running and try starting the application again. If SCA composite assets do not start, ensure that each asset is mapped to a deployment target that supports SCA composites.

If an asset composition unit uses an Enterprise JavaBeans (EJB) binding and does not start because it has a non-WebSphere target of "null", delete the asset composition unit and add it again to the business-level application. Specify a target that supports SCA composites when you add the asset to the business-level application. You cannot change the target after deployment.

If the SCA application uses security, the target server or cluster must be in the global security domain.

In multiple-node environments, synchronize the nodes after you save changes to the target before starting the business-level application.

Starting business-level applications

You can start a business-level application that is not running (has a status of Stopped). The application must contain code that can run on a server to start.

Before you begin

The application must be installed on a server. By default, the application starts automatically when the server starts.

About this task

You can start and stop business-level applications manually using the administrative console or wsadmin commands.

This topic describes how to use the administrative console to start a business-level application.

Procedure

1. Go to the Business-level applications page.
Click **Applications > Application Types > Business-level applications** in the console navigation tree.
2. Select the check box for the application you want started.
3. Click **Start**. The product runs the application and changes the state of the application to Started. The status is changed to partially started if not all servers on which the application is deployed are running.

Results

A message stating that the application started displays at the top the page.

If the business-level application does not start, ensure that the deployment target to which the application maps is running and try starting the application again.

If an application server on which the application is deployed synchronizes configuration with the deployment manager during server startup, then the application might not start and a `DeploymentDescriptorLoadException` error might be written to the server `SystemErr.log` file. Stop and restart the server, and then try starting the application again.

Note: This topic references one or more of the application server log files. Beginning in WebSphere Application Server Version 8.0 you can configure the server to use the High Performance Extensible Logging (HPEL) log and trace infrastructure instead of using `SystemOut.log`, `SystemErr.log`, `trace.log`, and `activity.log` files or native z/OS logging facilities. If you are using HPEL, you can access all of your log and trace information using the LogViewer command-line tool from your server profile bin directory. See the information about using HPEL to troubleshoot applications for more information on using HPEL.

If the application contains Service Component Architecture (SCA) composites and does not start, check for the following problems:

- If SCA composite assets do not start, ensure that each asset is mapped to a Version 8 deployment target or to a Version 7 deployment target that supports SCA composites.

- If an asset composition unit uses an Enterprise JavaBeans (EJB) binding and does not start because it has a non-WebSphere target of "null", delete the asset composition unit and add it again to the business-level application. Specify a target that supports SCA composites when you add the asset to the business-level application. You cannot change the target after deployment.
- If the META-INF/sca-deployables directory has multiple SCA composite files and the application does not start because the product cannot obtain the CompUnitInfoLoader value, place only the file that contains the composite in the META-INF/sca-deployables directory. You can place the other composite files anywhere else within the archive.

In multiple-node environments, synchronize the nodes after you save changes to the target before starting the business-level application.

What to do next

To restart a running application, select the application you want to restart, click **Stop** and then click **Start**.

Stopping business-level applications

You can stop a business-level application that is running and has a status of Started).

Before you begin

The application must be running on a product server.

About this task

You can stop applications manually using the administrative console or wsadmin commands.

This topic describes how to use the administrative console to stop a business-level application.

Procedure

1. Go to the Business-level applications page.
Click **Applications > Application Types > Business-level applications** in the console navigation tree.
2. Select the check box for the application you want stopped.
3. Click **Stop**. The product stops the processing of the application and changes the state of the application to Stopped.

Results

The status of the application changes and a message stating that the application stopped displays at the top the page.

What to do next

To restart a stopped application, select the application you want to restart, and then click **Start**.

Updating business-level applications

You can update business-level applications by deleting or changing composition units, or by mapping composition units to different deployment targets.

Before you begin

Determine the changes that you want to make to your application. Also, determine whether the changed application can run on your deployment targets.

The administrative console Server pages show the versions for deployment targets.

If you want to change a composition unit that contains an enterprise bundle archive (EBA) asset, see [Modifying the configuration of an EBA asset](#).

About this task

Updating consists of adding new composition units to an application, replacing or removing composition units, or mapping composition units to different deployment targets.

You can add an asset or shared library composition unit to multiple business-level applications. However, each composition unit for the same asset must have a unique composition unit name. You can add a business-level application composition unit to more than one business-level application.

This topic describes how to update business-level applications using the administrative console. Alternatively, you can use programming or the wsadmin tool.

Procedure

- Delete composition units from your business-level application.
 1. Go to the business-level application settings page.
 - Click **Applications** > **Application Types** > **Business-level applications** > *application_name* in the console navigation tree.
 2. Select each composition unit of the application that you want to delete.
 3. Click **Delete**.
 4. On the Delete composition unit from business-level application page, confirm the deletion and click **OK**.
- Add new or updated assets, shared libraries, or other business-level applications to your business-level application.
 1. Update asset binary files or shared libraries as needed.
 2. If you are adding new assets that are not registered with the product management domain, import the assets.
 3. If you are updating existing assets, use the **Update** option to update asset files.
 4. On the business-level application settings page, specify the type of composition unit to add.
 - To add an asset, under **Deployed assets**, click **Add** > **Add Asset**.
 - To add a shared library, under **Deployed assets**, click **Add** > **Add Shared Library**.
 - To add a business-level application, under **Business-level applications**, click **Add**.
 5. On the New composition unit page, select a unit from the list of available units, and then click **Continue**.
 6. On the Set options page, change the composition unit settings as needed, and then click **Next**.
 7. On the Map composition unit to a target page, change the deployment target as needed, and then click **Next**.
 - This page is not shown when you add a business-level application.
 8. On the Summary page, click **Finish**.
 9. If the product adds the unit successfully, click **Manage application**.
 - If the unit addition is not successful, read the messages, and try adding the unit again. Correct the errors noted in any messages.

10. On the Adding composition unit to the business-level application page, click **Save**.
11. Repeat these steps to add any other assets, shared libraries, or applications needed by the business-level application.

The business-level application settings page displays the configuration unit names.

- Map composition units to different deployment targets.
 1. On the composition unit settings page, select the composition unit that you want to change.
 2. Under **Current targets**, click **Modify Target**.
 3. On the Map targets page, change the target.
 - a. From the list of available clusters and servers, select a different deployment target.
 - b. Click >> to add the deployment target to the **Selected** list.
 - c. To remove a deployment target from the **Selected** list, select the target and click <<.
 - d. Click **OK**.

The business-level application settings page displays the selected deployment target.

What to do next

Save the changes to your administrative configuration.

When saving the configuration, synchronize the configuration with the nodes where the application is expected to run.

Updating SCA composite artifacts

You can view and update Service Component Architecture (SCA) composite artifacts in business-level applications.

Before you begin

Add an SCA artifact as a composition unit to a business-level application.

About this task

You can view and update the following SCA composite artifacts:

- Composite level property definition
- Composite level component property definition
- Composite level component reference definition

You can view and update SCA composite artifacts using the administrative console or the wsadmin tool. This topic describes how to view and update SCA composite artifacts using the administrative console.

Procedure

1. Go to the composition unit settings page for an SCA composite artifact in a business-level application.
Click **Applications > Application Types > Business-level applications > *application_name* > *SCA_deployed_asset_composition_unit_name***.

The composition unit settings page for an SCA composite artifact has fields that are not shown on the composition unit settings page for a non-SCA artifact:

- SCA composite components
 - SCA composite properties
 - SCA composite wires
2. Click on a name link in one of these SCA fields to view the settings for an SCA artifact.
The SCA fields display None instead of a name link if the composition unit does not have that particular type of SCA composite.

3. Optional: Update a SCA composite setting value.
 - a. Change an existing setting value for the SCA artifact.
 - b. Click **OK**.

The setting value is updated.

Viewing and updating SCA composites in HelloWorldAsync

“Example: Creating an SCA business-level application with the console” on page 455 describes how to create the HelloWorldAsync business-level application. This application contains an SCA artifact, helloworldws, as a composition unit. You can view and update settings for SCA composites in the helloworldws composition unit using the console.

1. Go to the composition unit settings page for the helloworldws composition unit in the HelloWorldAsync business-level application.

Click **Applications > Application Types > Business-level applications > HelloWorldAsync > helloworldws**.

From the composition unit settings page, you can view information associated with helloworldws, as well as update composite settings.

2. Click on a link for the SCA artifact to be viewed or updated.

For example, click on the HelloWorldServiceComponent link under **SCA composite components** and, in the page that displays, click on the HelloWorldService link under **Service**. In the Component service settings page that displays, you can specify a setting value for the service.

3. If you update a setting value for the SCA artifact, click **OK**.

What to do next

Save the changes to your administrative configuration.

On multiple-server products, when saving the configuration, synchronize the configuration with the nodes where the application is expected to run.

Viewing SCA composite definitions

You can view information on the definition of a Service Component Architecture (SCA) composite in the administrative console.

Before you begin

The SCA composite must be a composition unit in a business-level application.

About this task

The composite definition provides data on the composite, such as component names and service references. The View composite page displays the composite definition of an SCA deployed asset composition unit.

Procedure

1. Go to the View composite page.

Click **Applications > Application Types > Business-level applications > *application_name* > *SCA_deployed_asset_name* > View composite**.
2. Optional: Click **Expand All** or **Collapse All** to more easily browse the page.

Results

The View composite page displays the contents of the composition unit definition.

Example

Suppose the HelloWorldAsync business-level application provided as a sample with the product is installed. Click **Applications > Application Types > Business-level applications > HelloWorldAsync > helloworldws > View composite**.

The View composite page displays configuration information resembling the following:

```
<composite targetNamespace="http://helloworld" name="helloworldws" >
  <component name="AsynchTranslatorComponent" >
    <implementation.java class="helloworld.impl.AsynchTranslatorComponent" />
    <service name="AsynchTranslatorService">
      <interface.java interface="helloworld.AsynchTranslatorService"
        callbackInterface="helloworld.HelloWorldCallback" />
      <binding.ws/>
      <callback>
        <binding.ws/>
      </callback>
    </service>
  </component>
</composite>
```

What to do next

Browse the page to ensure that it contains the intended configuration information.

Viewing SCA domain information

You can view information on Service Component Architecture (SCA) composites in an SCA domain in the administrative console.

Before you begin

The SCA composite must be a composition unit in a business-level application.

About this task

Viewing SCA domain information enables you to see on one console page information on all components in an SCA domain. The View domain page displays information on available services in the current domain.

Procedure

1. Go to the View domain page.
Click **Applications > Application Types > Business-level applications > *application_name* > *SCA_deployed_asset_name* > View domain**.
2. Optional: Click **Expand All** or **Collapse All** to more easily browse the page.

Results

The View domain page lists information on components in the current domain.

Example

Suppose the HelloWorldAsync business-level application provided as a sample with the product is installed. Click **Applications > Application Types > Business-level applications > HelloWorldAsync > helloworldws > View domain**.

The View domain page displays information resembling the following:

```
<domain name="myCell02">
  <component name = "AsynchTranslatorComponent"
    mapTarget = "WebSphere:cell=myCell02,node=myNode02,server=server1">
    <service name = "AsynchTranslatorService">
      <interface.java interface = "helloworld.AsynchTranslatorService"/>
    </service>
    <reference name = "AsynchTranslatorService" target = ""/>
    <httpurlendpoints name = "endpoints" uri = ""/>
  </component>
</domain>
```

What to do next

Browse the page to ensure that it contains the intended information.

You can export the same domain information to a file using the `exportCompositeToDomain` command. See “Exporting SCA domain information using scripting.”

Viewing and editing JMS bindings on references and services of SCA composites

You can view information on a Java Message Service (JMS) binding for a Service Component Architecture (SCA) composite in the administrative console. The JMS bindings page of the console displays the settings of a `binding.jms` element. You can use the console page to edit resource and response resource settings.

Before you begin

Configure an SCA composite that uses JMS bindings and add that composite as a composition unit to a business-level application. To view and edit settings for a JMS binding on references, the composite must define a JMS binding in the reference. Similarly, to view and edit settings for a JMS binding on services, the composite must define a JMS binding in the service.

About this task

The JMS bindings page enables you to see JMS binding settings for an SCA composite in the administrative console. The information shown is similar to that shown by running the `viewCompUnit wsadmin` scripting command to view the SCA composite composition unit. However, the console page shows JMS binding settings only, and does not show other information on the composition unit that running `viewCompUnit` returns.

After deployment of an SCA composite that uses JMS bindings, you can edit JMS binding resource or response resource settings that specify Java Naming and Directory Interface (JNDI) names on the JMS bindings page. The editing capabilities are similar to those of the `editCompUnit` scripting command. To edit settings, the JMS resource must exist. The product does not dynamically create JMS resources when you edit a composition unit.

Procedure

1. Go to the JMS bindings page.
To view this page, your composition unit must support a JMS binding reference or service.
 - a. Click **Applications > Application Types > Business-level applications > *application_name* > *deployed_asset_composition_unit_name***.
 - b. From the composition unit settings page for SCA composites, select to view references or services:
 - For an SCA component reference, click **SCA Composite Components > *SCA_component_reference_name* > SCA Component References > *reference_name* > Bindings > JMS binding**.
 - For an SCA component service, click **SCA Composite Components > *SCA_component_service_name* > SCA Component Services > *service_name* > Bindings > JMS binding**.
2. View the JMS binding settings for the composition unit.
3. Edit the **Destination JNDI name**, **Activation specification JNDI name**, **Connection factory JNDI name**, **Response destination JNDI name**, or **Response connection factory JNDI name** settings as needed and click **Apply**.

Results

The JMS bindings page displays property settings for a JMS binding. If you change resource or response resource JNDI name values, the changed values are shown.

Example

Suppose that you have a business-level application named MyJmsBLA, which has an SCA composite named MySCAComposite. This composite uses a JMS binding on service named JmsService, in a composition unit named myJmsBindingCU.

1. Click **Applications > Application Types > Business-level applications > MyJmsBLA > myJmsBindingCU > SCA Composite Components > MySCAComposite > SCA Component Services > JmsService > Bindings > JMS binding**.
The JMS bindings page displays general binding settings and settings for a JMS binding on services.
2. Edit the value specified for **Activation specification JNDI name**, and click **Apply**.
The JMS bindings page displays the changed value.

What to do next

Browse the page to ensure that it contains the binding information.

For information on JMS binding settings, refer to the online help for the JMS bindings page or Section 1.4 of the SCA JMS Binding specification, Version 1.00.

Exporting WSDL and XSD documents

You can export Web Services Description Language (WSDL) and XML schema definition (XSD) documents used by a Service Component Architecture (SCA) composition unit to a location of your choice.

Before you begin

Your SCA business-level application must contain one or more composition units that use a WSDL or XSD document.

A WSDL document is a file that provides a set of definitions that describe a web service in WSDL, an Extensible Markup Language (XML)-based description language.

An XSD document is an instance of an XML schema written in the XML schema definition language. The document has the extension `.xsd`. The prefix `xsd` in the XML elements of an XSD document indicates the XML schema namespace.

About this task

Note: You can export WSDL and XSD documents that are used by an SCA composition unit using the administrative console. In previous releases, you had to use the `exportWSDLArtifacts` command to export WSDL and XSD documents.

On the composition unit settings page for an SCA composite, click the **Export WSDL and XSD documents** link and then specify the target location for the files.

The product extracts from the selected composition unit the WSDL and XSD files that are required for web services client development. The files are for the services exposed by the web service binding, `binding.ws`.

Procedure

1. Go to the composition unit settings page for the SCA composite.
Click **Applications > Application Types > Business-level applications > *application_name* > *SCA_deployed_asset_name***.
2. Click the **Export WSDL and XSD documents** link.
3. From the displayed dialog, specify the target directory to which to save the documents.

Results

The WSDL and XSD documents are copied to the target directory.

Example

Suppose you want to export WSDL or XSD documents in the `HelloWorldAsync` business-level application that the product provides as a sample. Complete the following actions in the administrative console:

1. Click **Applications > Application Types > Business-level applications > HelloWorldAsync > helloworldws > Export WSDL and XSD documents**.
2. Using the displayed dialog, specify a directory that exists on your computer.

The product adds the `helloworldws_WSDLArtifacts.zip` file to the specified directory. The `helloworldws_WSDLArtifacts.zip` file has one WSDL file, `AsynchTranslatorComponentTranslatorService_wsdlgen.wsdl`.

What to do next

Examine the exported files to ensure that they contain the intended WSDL and XSD documents.

You can export WSDL and XSD documents using the `exportWSDLArtifacts` command. See "Exporting WSDL and XSD documents using scripting."

Deleting business-level applications

After an application no longer is needed, you can delete it.

About this task

Deleting a business-level application removes the application from the product configuration repository and it deletes the application binaries from the file system of all nodes where the application files are installed.

Procedure

1. Go to the Business-level applications page.
Click **Applications > Application Types > Business-level applications** in the console navigation tree.
2. If you need to retain a copy of the application, back up composition units of the application.
3. Delete composition units of the application.
 - a. On the Business-level applications page, click the name of the business-level application that you want to delete.
 - b. On the business-level application settings page, delete each composition unit of the application. Deployed assets and business-level applications can be composition units of a business-level application.
Select one or more composition units and click **Delete**.
For Service Component Architecture (SCA) business-level applications that use an OSGi implementation in an enterprise bundle archive (EBA) composition unit, delete the SCA composition unit before attempting to delete the EBA composition unit, otherwise the product returns an error. An EBA composition unit that is used for the implementation cannot be deleted until the SCA composition unit that uses the `implementation.osgiapp` is deleted.
 - c. On the Delete composition unit from Business-level application page, confirm the deletion and click **OK**.
 - d. Repeat steps b and c until the business-level application that you want to delete has no more composition units.

Deleting a composition unit removes the configuration from the `profile_root/config/cells/cell_name/cus` directory.

4. Delete the business-level application.
 - a. Select the application that you want to delete.
 - b. Click **Delete**.

Unless the application is used by another business-level application, deleting a business-level application removes the configuration from the `profile_root/config/cells/cell_name/bias` directory.

5. On the Delete business-level application page, confirm the deletion and click **OK**.
6. Save changes made to the administrative configuration.

Results

On single-server products, application binaries are deleted after you save the changes.

On multiple-server products, application binaries are deleted when configuration changes on the deployment manager synchronize with configurations for individual nodes.

Deleting the HelloWorldAsync business-level application

“Example: Creating an SCA business-level application with the console” on page 455 describes how to create the HelloWorldAsync business-level application. You can delete this application using the console.

1. Go to the Business-level applications page and, if HelloWorldAsync is running, change its status to Stopped.
 - a. Click **Applications > Application Types > Business-level applications**.
 - b. Select HelloWorldAsync.

- c. Click **Stop**.
2. Go to the business-level applications settings page for HelloWorldAsync and delete the helloworldws composition unit.
 - a. Click **Applications > Application Types > Business-level applications > HelloWorldAsync**.
 - b. From **Deployed assets**, select helloworldws.
 - c. Click **Delete**.
 - d. On the Delete composition unit from Business-level application page, confirm the deletion and click **OK**.
 - e. Click the **Save** link to save the changes.
3. From the business-level applications page, delete the HelloWorldAsync application.
 - a. Click **Applications > Application Types > Business-level applications**.
 - b. Select HelloWorldAsync.
 - c. Click **Delete**.
 - d. On the Delete business-level application page, click **OK**.
 - e. Click the **Save** link to save the changes.
4. Optionally, from the Assets page, delete the helloworld-ws-async.jar asset from the asset repository.
 - a. Click **Applications > Application Types > Assets**.
 - b. Select helloworld-ws-async.jar.
 - c. Click **Delete**.
 - d. On the Delete asset page, click **OK**.
 - e. Click the **Save** link to save the changes.

What to do next

If using the administrative console **Delete** options does not fully delete a business-level application or its composition units, you can delete the business-level application and its composition units manually from a deployment manager or stand-alone server. Suppose you want to delete a business-level application named ExampleBLA, and ExampleBLA is not used by another business-level application. Complete the following steps to manually delete the ExampleBLA configurations from the blas and cus directories:

1. Delete the *profile_root/config/cells/cell_name/blas/ExampleBLA* directory.
2. Delete the *profile_root/config/cells/cell_name/cus/ExampleBLA* directory.
3. Save changes made to the administrative configuration.
4. On multiple-server products, synchronize the deployment manager with node configurations.

Chapter 12. Administering business-level applications using programming

You can use the command framework programming to create, edit, update, start, stop, delete, export, import, and query information about business-level applications. A business-level application defines an enterprise-level application.

Before you begin

This task assumes a basic familiarity with the command framework. Read about the command framework in the application programming interfaces documentation.

About this task

Besides creating, editing, updating, starting, stopping, deleting, exporting, importing, and querying information about business-level applications using programming, you can do these tasks using the administrative console or the wsadmin scripting tool.

Procedure

1. Perform any of the following tasks to administer your business-level applications using programming.
 - a. Create an empty business-level application.

You typically create an empty business-level application and then add assets or business-level applications as composition units to the empty business-level application.
 - b. Import an asset.

You can import an asset to register the asset with the product and optionally store the asset in the product repository so that you can later use the asset in a business-level application. An asset represents at least one binary file that implements business logic.
 - c. Add a composition unit.

You can add an asset to a business-level application by creating a composition unit for the asset. A composition unit is typically created from an asset and contains configuration information that makes the asset runnable.
 - d. Start a business-level application.

You can start a business-level application, which starts each composition unit in that business-level application. Each composition unit is started on the respective targets on which the business-level application is deployed.
 - e. Stop a business-level application.

You can stop a business-level application, which stops each composition unit in that business-level application. Each composition unit is stopped on the respective targets on which the business-level application is deployed.
 - f. Check the status of a business-level application.

You can check the status of an entire business-level application. You can also limit the status to a particular composition unit of a business-level application, a specific deployment target, or check the status of the composition unit and the deployment target at the same time.
 - g. Delete a business-level application.

You can delete a business-level application using programming. You might delete a business-level application if the application is not functioning correctly, no longer needed, and so on.
 - h. Delete an asset.

You can delete an asset from a business-level application using programming if the asset is not functioning correctly, the asset is no longer needed, and so on. An asset represents at least one binary file that implements business logic.

- i. Delete a composition unit.

You can delete a composition unit from a business-level application if the composition unit is not functioning correctly, the composition unit is no longer needed, and so on. A composition unit is typically created from a business-level application or an asset and contains configuration information that makes the asset runnable.
- j. Export an asset.

You can export an asset from the current session so that you can back up the asset, import the asset to another session, and so on. An asset represents at least one binary file that implements business logic.
- k. List assets.

You can list the assets that have been imported to the current workspace so that you can do further asset administration, such as deleting or exporting assets. An asset represents at least one binary file that implements business logic.
- l. List composition units.

You can list the composition units for a specific business-level application in a session so that you can do further composition unit administration, such as deleting or adding composition units. A composition unit is typically created from a business-level application or an asset and contains configuration information that makes the asset runnable.
- m. List business-level applications.

You can list the business-level applications of a session so that you can do further business-level application administration such as deleting a business-level application. A business-level application is an administrative model that captures the definition of an enterprise-level application so that you can perform specific business functions, such as accounting.
- n. Edit a composition unit.

You can edit the configuration information in a composition unit of a business-level application if, for example, you want to change which modules in the composition unit are configured to run in which targets. A composition unit is typically created from a business-level application or an asset and contains configuration information that makes the asset runnable.
- o. Edit an asset.

You can edit the information of an asset, for example, its destination location, its relationship with other assets, and so on. An asset represents at least one binary file that implements business logic
- p. Edit a business-level application.

You can edit the information of a business-level application such as its description. A business-level application is an administrative model that captures the entire definition of an enterprise-level application.
- q. Update an asset.

You can update an asset by adding, deleting, or updating a single file or Java Platform, Enterprise Edition (Java EE) module, or by merging multiple files or Java EE modules into an asset. You can also update an asset by replacing the entire asset.
- r. View a composition unit.

You can view the composition unit information so that you can do other tasks associated with the composition unit, such as editing an asset or deleting a composition unit. A composition unit is typically created from a business-level application or an asset and contains configuration information that makes the asset runnable.
- s. View an asset.

You can view the asset information so that you can do other tasks associated with the asset, such as editing or exporting an asset. An asset represents at least one binary file that implements business logic.
- t. View a business-level application.

You can view business-level application information such as the description so that you can do other tasks associated with the business-level application, such as editing the business-level application. A business-level application is an administrative model that captures the entire definition of an enterprise-level application.

u. List control operations.

You can list the control operations of a business-level application or a composition unit for a session. You use control operations, such as start or stop, to change or query the runtime environment of a business-level application or a composition unit.

2. Save your changes to the master configuration repository.
3. Synchronize changes to the master configuration across the nodes for the changes to take effect.

Results

Depending on which tasks you complete, you have created, edited, updated, started, stopped, deleted, exported, imported, or queried information about business-level applications.

What to do next

If you have further business-level application updates, you can do the updates through programming, the administrative console, or the wsadmin scripting tool.

Creating an empty business-level application using programming

You can create an empty business-level application, and then add assets or business-level applications as composition units to the empty business-level application.

Before you begin

This task assumes a basic familiarity with command framework programming. Read about command framework programming in the application programming interfaces documentation.

You can create an empty business-level application using programming, the administrative console, or the wsadmin tool.

About this task

Perform the following steps to create an empty business-level application using programming. In your code that creates the empty business-level application, you must provide the name parameter. The name parameter specifies the name of the business-level application that you create.

Procedure

1. Connect to the application server.
The command framework allows the administrative command to be created and run with or without being connected to the application server. This step is optional if the application server is not running.
2. Create the command manager.
The command manager provides the functionality to create a new administrative command or query existing administrative commands.
3. Optionally create the asynchronous command handler for listening to command notifications.
Business-level application commands are implemented as asynchronous commands. To monitor the progress of the running command, you have to create an asynchronous command handler to receive notifications that the command generates.
4. Create the asynchronous command client.

An asynchronous command client provides a higher level interface to work with an asynchronous command. If you created an asynchronous command handler in the previous step, the handler is passed to the asynchronous command client. The asynchronous command client forwards the command notification to the handler and helps to control running of the command.

5. Use the command manager that you created in a previous step to create and set up the command that creates an empty business-level application.

The command name is `createEmptyBLA`. The name parameter is a required parameter that you use to specify the name of the business-level application. You can optionally provide the description parameter to provide a description of the newly created business-level application.

6. Call the `processCommandParameters` method in the asynchronous command client to process the command parameters.

The command framework asynchronous command model requires this call.

7. Call the `execute` method in the asynchronous command client to run the command that creates an empty business-level application.

You might have created an asynchronous command handler to implement the `AsyncCommandHandlerIF` interface class in a previous step. If you did, the asynchronous command client listens to command notifications and forwards the notifications to the handler. The handler performs any necessary actions while waiting for the command to complete.

8. Check the command result when the command completes.

When the command finishes running, control is returned to the caller. You can then check the result by calling the `command.getCommandResult` method.

Results

After you successfully run the code, the empty business-level application is created.

Example

The following example shows how to create an empty business-level application based on the previous steps. Some statements are split on multiple lines for printing purposes.

```
package com.ibm.ws.management.application.task;

import java.util.Properties;

import com.ibm.websphere.management.AdminClient;
import com.ibm.websphere.management.AdminClientFactory;
import com.ibm.websphere.management.Session;
import com.ibm.websphere.management.cmdframework.AdminCommand;
import com.ibm.websphere.management.cmdframework.CommandMgr;
import com.ibm.websphere.management.cmdframework.CommandResult;
import com.ibm.websphere.management.exception.AdminException;
import com.ibm.websphere.management.async.client.AsyncCommandClient;

public class CreateEmptyBLA {

    public static void main(String[] args) {

        try {

            // Connect to the application server.
            // This step is optional if you use the local
            // command manager. Comment out the lines to and including
            // CommandMgr cmdMgr = CommandMgr.getClientCommandMgr(
            // soapClient);
            // to get the soapClient soap client if you use the local
            // command manager.

            String host = "localhost";
```

```

String port = "8880";
// Change to your port number if it is
// not 8880.

Properties config = new Properties();
config.put(AdminClient.CONNECTOR_HOST, host);
config.put(AdminClient.CONNECTOR_PORT, port);
config.put(AdminClient.CONNECTOR_TYPE,
    AdminClient.CONNECTOR_TYPE_SOAP);
System.out.println("Config: " + config);
AdminClient soapClient =
    AdminClientFactory.createAdminClient(config);

// Create the command manager.
CommandMgr cmdMgr = CommandMgr.getClientCommandMgr(soapClient);

// Comment out the previous lines to create a client command
// manager if you are using a local command manager.
// Uncomment the following line to create a local command
// manager:
//
// CommandMgr cmdMgr = CommandMgr.getCommandMgr();
System.out.println("\nCreated command manager");

// Optionally create an asynchronous command handler
// for listening to command notifications.
// Comment out the following line if no further handling
// of command notification is required:
AsyncCmdTaskHandler listener = new AsyncCmdTaskHandler();

// Create an asynchronous command client.

// Set up the session.
String id = Long.toHexString(System.currentTimeMillis());
String user = "content" + id;
Session session = new Session(user, true);

// If no command handler is used, replace the listener with
// null for the AsyncCommandClient object that follows.

AsyncCommandClient asyncCmdClientHelper = new
    AsyncCommandClient(session, listener);
System.out.println("\nCreated async command client");

// Create the command that creates an empty
// business-level application.
String cmdName = "createEmptyBLA";

AdminCommand cmd = cmdMgr.createCommand(cmdName);
cmd.setConfigSession(session); // Create an empty
    // business-level application using
    // the session created.
System.out.println("\nCreated " + cmdName);

// Set the name command parameter.
String blaName = "bla1";
cmd.setParameter("name", blaName);

System.out.println("\nSet name parameter to "
    + cmd.getParameter("name"));

// Uncomment the following lines to set the description of
// the business-level application being created:
//
// String blaDescription = "description for bla1";
// cmd.setParameter("description", blaDescription);
// System.out.println("\nSet description parameter to " +

```

```

//          cmd.getParameter("description"));

// Call the asynchronous command client to
// process the command parameters.
try {
    asyncCmdClientHelper.processCommandParameters(cmd);
    System.out.println("\nCompleted process command " +
        "parameters");
} catch (Throwable th) {
    System.out.println("Failed from " +
        "asyncCmdClientHelper.processCommandParameters(cmd).");
    th.printStackTrace();
    System.exit(-1);
}

// Call the asynchronous command client to run the command.
asyncCmdClientHelper.execute(cmd);
System.out.println("\nCompleted command execution");

// Check the command result.
CommandResult result = cmd.getCommandResult();
if (result != null) {
    if (result.isSuccessful()) {
        System.out.println("\nCommand ran successfully "
            + "with result\n" + result.getResult());
    } else {
        System.out.println("\nCommand ran with " +
            "Exception");
        result.getException().printStackTrace();
    }
}
} catch (Exception e) {
    e.printStackTrace();
}
}

package com.ibm.ws.management.application.task;

import com.ibm.websphere.management.cmdframework.provider.CommandNotification;
import com.ibm.websphere.management.async.client.AsyncCommandHandlerIF;

public class AsyncCmdTaskHandler implements AsyncCommandHandlerIF {

    public void handleNotification(CommandNotification notification) {
        // Add your own code here to handle the received notification.
        System.out.println("\nEXAMPLE: notification received: " +
            notification);
    }
}

```

What to do next

You can add business-level applications or assets as composition units into the newly created business-level application. Alternatively, you can add the newly created business-level application to other business-level applications.

Importing an asset using programming

You can import an asset to register the asset with the product and optionally store the asset in the product repository so that you can later use that asset in a business-level application. An asset represents at least one binary file that implements business logic.

Before you begin

This task assumes a basic familiarity with command framework programming. Read about command framework programming in the application programming interfaces documentation.

You can import an asset using programming, the administrative console, or the wsadmin tool.

About this task

When you import an asset, you register the asset with the product and optionally store the asset in the product repository.

You must provide a file path to the source that you are importing. Specify an absolute path name to the source, as the behavior for a relative path is unpredictable.

You can specify a destination location from where the application server reads the asset file while starting a composition unit created from the asset. The asset is copied to this location when the configuration changes are synchronized with the node. The default asset destination is *profile_root/installedAssets/asset_name*.

You can optionally specify a storage type of FULL, METADATA, or NONE. The default value is FULL, which means that the asset and associated meta data are stored in the product asset repository. If you specify a storage type of METADATA, the asset is not copied to the product repository, but associated meta data is stored in the product repository. If you specify a storage type of NONE, neither the asset nor the asset meta data is stored in the product asset repository. For storage types of METADATA and NONE, the asset is expected to reside at the destination file path. Storage types of METADATA and NONE are typically used by development tools which enable iterative development on the copy of the asset in the directory structure of the tool.

Perform the following steps to import an asset using programming.

Procedure

1. Connect to the application server.

The command framework allows the administrative command to be created and run with or without being connected to the application server. This step is optional if the application server is not running.

2. Create the command manager.

The command manager provides the functionality to create a new administrative command or query existing administrative commands.

3. Optionally create the asynchronous command handler for listening to command notifications.

Business-level application commands are implemented as asynchronous commands. To monitor the progress of the running command, you have to create an asynchronous command handler to receive notifications that the command generates.

4. Create the asynchronous command client.

An asynchronous command client provides a higher level interface to work with an asynchronous command. If you created an asynchronous command handler in the previous step, the handler is passed to the asynchronous command client. The asynchronous command client forwards the command notification to the handler and helps to control running of the command.

5. Use the command manager that you created in a previous step to set up the command that imports an asset.

The command name is `importAsset`. The source parameter is a required parameter that you use to specify the path to the asset. You can optionally provide the `storageType` parameter to specify how to save the asset in the configuration repository.

6. Call the `processCommandParameters` method in the asynchronous command client to process the command parameters.

The command framework asynchronous command model requires this call.

7. Set up the command step parameters.

You can set parameters in the `AssetOptions` step that contains data about the asset such as its description, file permission, and relationship with other assets.

8. Call the asynchronous command client to run the command that imports an asset.

You might have created an asynchronous command handler to implement the `AsyncCommandHandlerIF` interface class in a previous step. If you did, the asynchronous command client listens to command notifications and forwards the notifications to the handler. The handler performs any necessary actions while waiting for the command to complete.

9. Check the command result when the command completes.

When the command finishes running, control is returned to the caller. You can then check the result by calling the `command.getCommandResult` method.

Results

After you successfully run the code, the asset is imported.

Example

The following example shows how to import an asset based on the previous steps.

Some statements are split on multiple lines for printing purposes.

```
package com.ibm.ws.management.application.task;

import java.util.Properties;
import com.ibm.websphere.management.AdminClient;
import com.ibm.websphere.management.AdminClientFactory;
import com.ibm.websphere.management.Session;

import com.ibm.websphere.management.cmdframework.CommandMgr;
import com.ibm.websphere.management.cmdframework.AdminCommand;
import com.ibm.websphere.management.cmdframework.TaskCommand;
import com.ibm.websphere.management.cmdframework.CommandStep;
import com.ibm.websphere.management.cmdframework.CommandResult;
import com.ibm.websphere.management.cmdframework.UploadFile;
import com.ibm.websphere.management.exception.AdminException;
import com.ibm.websphere.management.async.client.AsyncCommandClient;

public class ImportAsset {

    public static void main (String [] args) {

        try {

            // Connect to the application server.
            // This step is optional if you use the local
            // command manager. Comment out the lines to and including
            // CommandMgr cmdMgr = CommandMgr.getClientCommandMgr(
            // soapClient);
            // to get the soapClient soap client if you use the local
            // command manager.

            String host = "localhost";
            String port = "8880"; //Change to your port number if it is not
            //8880.

            Properties config = new Properties();
            config.put (AdminClient.CONNECTOR_HOST, host);
```



```

config.put (AdminClient.CONNECTOR_PORT, port);
config.put (AdminClient.CONNECTOR_TYPE,
            AdminClient.CONNECTOR_TYPE_SOAP);
System.out.println ("Config: " + config);
AdminClient soapClient =
            AdminClientFactory.createAdminClient(config);

// Create the command manager.
CommandMgr cmdMgr = CommandMgr.getClientCommandMgr(soapClient);

// Comment out the previous lines to create a client command
// manager if you are using a local command manager.
// Uncomment the following line to create a local command
// manager:
//
//CommandMgr cmdMgr = CommandMgr.getCommandMgr();

System.out.println("\nCreated command manager");

// Optionally create the asynchronous command handler.
// Comment out the following line if no further handling
// of command notification is required:
AsyncCmdTaskHandler listener = new AsyncCmdTaskHandler();

// Create an asynchronous command client.

// Setup the session.
String id = Long.toHexString(System.currentTimeMillis());
String user = "content" + id;
Session session = new Session(user, true);

// If no command handler is used, replace the following listener with
// null for the AsyncCommandClient object.

AsyncCommandClient asyncCmdClientHelper = new
            AsyncCommandClient(session, listener);
System.out.println("\nCreated async command client");

String cmdName = "importAsset";
UploadFile assetSource =
UploadFile("/sources/test5.zip"); //Change to the directory of your sources.

// Create the command to import an asset.
AdminCommand cmd = cmdMgr.createCommand(cmdName);
cmd.setConfigSession(session); //import the asset using
//the session created
System.out.println("\nCreated " + cmdName);

// Set the source command parameter.
cmd.setParameter("source", assetSource);
System.out.println("\nSet source parameter to " +
            cmd.getParameter("source"));

// Uncomment the following line to set the storage type to
// a value of STORAGETYPE_META or STORAGETYPE_NONE instead of
// the default of STORAGETYPE_FULL:
//
//cmd.setParameter("storageType,
//            CommandConstants.STORAGETYPE_NONE);

// Call the asynchronous client helper to process parameters.
try {
    asyncCmdClientHelper.processCommandParameters(cmd);
    System.out.println("\nCompleted process command " +
            "parameters");
}

```

```

catch(Throwable th) {
    System.out.println("Failed from " +
        "asyncCmdClientHelper.processCommandParameters(cmd).");
    th.printStackTrace();
    System.exit(-1);
}

// Set up the step parameters for the AssetOptions step.
String stepName = "AssetOptions";
CommandStep step = ((TaskCommand) cmd).gotoStep(stepName);

// The new asset name must contain the
// same extension as the original .zip file name.
String assetNewName = "asset1.zip";

// If you override the default destination, include the
// entire path with the file name for the new destination.
String destName = "/websphere/asset/installDir/asset1.zip";
for (int i = 0; i < step.getNumberOfRows(); i++) {
    // The following lines change the name and destination
    // step parameters. Other step parameters that you can
    // use follow, but are commented out.
    // Change your set
    // of step parameters as required by your scenario.

    // Set the name.
    step.setParameter("name", assetNewName, i);
    System.out.println("\nSet name parameter to " +
        step.getParameter("name", i));

    // Set the destination.
    step.setParameter("destination", destName, i);
    System.out.println("\nSet destination parameter to " +
        step.getParameter("destination", i));

    // Set the description.
    //String desc = "description for asset1.zip";
    //step.setParameter("description", desc, i);
    //System.out.println("\nSet description parameter to " +
    //    step.getParameter("description", i));

    // Set the validation.
    //String validate = "Yes";
    //step.setParameter("validate", validate, i);
    //System.out.println("\nSet validate parameter to " +
    //    step.getParameter("validate", i));

    // Set the file permission.
    //String filePermission = ".*\\.dll=755";
    //step.setParameter("filePermission", filePermission, i);
    //System.out.println("\nSet filePermission parameter to " +
    //    step.getParameter("filePermission", i));

    // Set the type aspect parameter value.
    // Format for a typeAspect: WebSphere:spec=xxx,version=n.n+
    // Websphere:spec=xxx,version=n.n.
    //String typeAspect = "";
    //step.setParameter("typeAspect", typeAspect, i);
    //System.out.println("\nGet typeAspect: " +
    //    step.getParameter("typeAspect", i));

    // Set the relationship parameter.
    // The relationship parameter declares dependency
    // relationships on other assets. The parameter value
    // is a list which contains the ID of each asset declared
    // as a dependency. Each ID in the list is separated by

```

```

        // a "plus" sign ("+").
        //
        // Only assets which are Java archives can be referenced in
        // dependency relationships. An asset is a Java archive if
        // it has a type aspect identifying it as such.
        //
        // If an asset declared as a dependency does not exist or
        // does not have a Java archive type aspect, it is ignored
        // and no dependency on the asset is registered in the
        // asset's configuration.
        //
        //String relationship =
        //    "assetname=shared.zip+assetname=shared2.zip";
        //step.setParameter("relationship", relationship, i);
        //System.out.println("\nGet relationship: " +
        //    step.getParameter("relationship", i));
    }

    // Call the asynchronous command client that imports the asset.
    asyncCmdClientHelper.execute(cmd);
    System.out.println("\nCompleted running of command");

    // Check the command result.
    CommandResult result = cmd.getCommandResult();
    if (result != null) {
        if (result.isSuccessful()) {
            System.out.println("\nCommand ran successfully " +
                "with result\n" + result.getResult());
        } else {
            System.out.println("\nCommand ran with " +
                "exception");
            result.getException().printStackTrace();
        }
    }
} catch (Exception e) {
    e.printStackTrace();
}
}
}

package com.ibm.ws.management.application.task;

import com.ibm.websphere.management.cmdframework.provider.CommandNotification;
import com.ibm.websphere.management.async.client.AsyncCommandHandlerIF;

public class AsyncCmdTaskHandler implements AsyncCommandHandlerIF {

    public void handleNotification(CommandNotification notification) {
        // Add your own code here to handle the received notification
        System.out.println("\nEXAMPLE: notification received: " +
            notification);
    }
}
}

```

What to do next

Add a composition unit to a business-level application using the asset that you imported. An asset included in a business-level application is represented by a composition unit.

Listing assets using programming

You can list the assets that have been imported so that you can do further asset administration, such as deleting or exporting assets. An asset represents at least one binary file that implements business logic.

Before you begin

This task assumes a basic familiarity with command framework programming. Read about command framework programming in the application programming interfaces documentation.

You can list assets using programming, the administrative console, or the wsadmin tool.

About this task

You can list assets using programming, the administrative console, or the wsadmin tool. This topic describes how to list assets using programming.

When you list assets, all the assets are listed unless you set the assetID to specify the asset that you want to list. You can optionally include deployable units or a description of the assets when you list the assets. After you list the assets, you can use the information to do further administration, such as deleting or exporting assets.

Perform the following tasks to list assets using programming.

Procedure

1. Connect to the application server.
The command framework allows the administrative command to be created and run with or without being connected to the application server. This step is optional if the application server is not running.
2. Create the command manager.
The command manager provides the functionality to create a new administrative command or query existing administrative commands.
3. Optionally create the asynchronous command handler for listening to command notifications.
Business-level application commands are implemented as asynchronous commands. To monitor the progress of the running command, you have to create an asynchronous command handler to receive notifications that the command generates.
4. Create the asynchronous command client.
An asynchronous command client provides a higher level interface to work with an asynchronous command. If you created an asynchronous command handler in the previous step, the handler is passed to the asynchronous command client. The asynchronous command client forwards the command notification to the handler and helps to control running of the command.
5. Use the command manager that you created in a previous step to create and set up the command that lists assets.
The command name is listAssets. You can optionally set the assetID parameter to query for assets that match the ID. You can also optionally set the includeDescription parameter and the includeDeplUnit parameter to include the display of the asset description and its deployable units.
6. Call the processCommandParameters method in the asynchronous command client to process the command parameters.
The command framework asynchronous command model requires this call.
7. Call the asynchronous command client to list the asset.
You might have created an asynchronous command handler to implement the AsyncCommandHandlerIF interface class in a previous step. If you did, the asynchronous command client listens to command notifications and forwards the notifications to the handler. The handler performs any necessary actions while waiting for the command to complete.
8. Check the command result when the command completes.
When the command finishes running, control is returned to the caller. You can then check the result by calling the command.getCommandResult method.

Results

After you successfully run the code, a list of assets is displayed.

Example

The following example shows how to list the assets based on the previous steps. Some statements are split on multiple lines for printing purposes.

```
package com.ibm.ws.management.application.task;

import java.util.Properties;

import com.ibm.websphere.management.AdminClient;
import com.ibm.websphere.management.AdminClientFactory;
import com.ibm.websphere.management.Session;
import com.ibm.websphere.management.cmdframework.AdminCommand;
import com.ibm.websphere.management.cmdframework.CommandMgr;
import com.ibm.websphere.management.cmdframework.CommandResult;
import com.ibm.websphere.management.async.client.AsyncCommandClient;

public class ListAssets {

    public static void main(String[] args) {

        try {

            // Connect to the application server.
            // This step is optional if you use the local
            // command manager. Comment out the lines to and including
            // CommandMgr cmdMgr = CommandMgr.getClientCommandMgr(
            // soapClient);
            // to get the soapClient soap client if you use the local
            // command manager.

            String host = "localhost";
            String port = "8880"; // Change to your port number if
                // it is not 8880.

            Properties config = new Properties();
            config.put(AdminClient.CONNECTOR_HOST, host);
            config.put(AdminClient.CONNECTOR_PORT, port);
            config.put(AdminClient.CONNECTOR_TYPE,
                AdminClient.CONNECTOR_TYPE_SOAP);
            System.out.println("Config: " + config);
            AdminClient soapClient =
                AdminClientFactory.createAdminClient(config);

            // Create the command manager.
            CommandMgr cmdMgr = CommandMgr.getClientCommandMgr(soapClient);

            // Comment out the previous lines to create a client command
            // manager if you are using a local command manager.
            // Uncomment the following line to create a local command
            // manager:
            //
            // CommandMgr cmdMgr = CommandMgr.getCommandMgr();
            System.out.println("\nCreated command manager");

            // Optionally create an asynchronous command handler.
            // Comment out the following line if no further handling
            // of command notification is required.
            AsyncCmdTaskHandler listener = new AsyncCmdTaskHandler();

            // Create an asynchronous command client.
```

```

// Set up the session.
String id = Long.toHexString(System.currentTimeMillis());
String user = "content" + id;
Session session = new Session(user, true);

// If no command handler is used, replace listener with
// null for the AsyncCommandClient object.
AsyncCommandClient asyncCmdClientHelper = new
    AsyncCommandClient(session, listener);
System.out.println("\nCreated async command client");

// Create the command that lists the assets.
String cmdName = "listAssets";
AdminCommand cmd = cmdMgr.createCommand(cmdName);
cmd.setConfigSession(session); // list all the assets
                                // using the session created.
System.out.println("\nCreated " + cmdName);

// Optionally set the assetID parameter.
// Uncomment the following code to set the assetID parameter to
// only list the asset with the ID specified, otherwise all
// assets are listed. Change the assetID parameter according to your
// scenario.
// Examples of valid formats for the assetID parameter are:
// - aName
// - assetname=aName
// - WebSphere:assetname=aName
// All assets that match the ID specification are listed.
// The ID must include at least the asset name.
// String assetID = "asset1.zip";
// cmd.setParameter("assetID", assetID);

//System.out.println("\nSet assetID parameter to "
//    + cmd.getParameter("assetID"));

// Optionally include a description by setting
// the includeDescription parameter to true or false.
String includeDescription = "true";
cmd.setParameter("includeDescription", includeDescription);

System.out.println("\nSet includeDescription parameter to "
    + cmd.getParameter("includeDescription"));

// Optionally include deployable units by setting
// the includeDeplUnit parameter to true or false.
String includeDeplUnit = "false";
cmd.setParameter("includeDeplUnit", includeDeplUnit);

System.out.println("\nSet includeDeplUnit parameter to "
    + cmd.getParameter("includeDeplUnit"));

// Call the asynchronous client helper to process parameters.
try {
    asyncCmdClientHelper.processCommandParameters(cmd);
    System.out.println("\nCompleted process command " +
        "parameters");
} catch (Throwable th) {
    System.out.println("Failed from " +
        "asyncCmdClientHelper.processCommandParameters(cmd).");
    th.printStackTrace();
    System.exit(-1);
}

// Run the command to list assets.
asyncCmdClientHelper.execute(cmd);
System.out.println("\nCompleted running of the command");

```

```

// Check the command result.
CommandResult result = cmd.getCommandResult();
if (result != null) {
    if (result.isSuccessfull()) {
        System.out.println("\nCommand ran successfully "
            + "with result\n" + result.getResult());
    }
    else {
        System.out.println("\nCommand ran with " +
            "Exception");
        result.getException().printStackTrace();
    }
}
} catch (Exception e) {
    e.printStackTrace();
}
}

package com.ibm.ws.management.application.task;

import com.ibm.websphere.management.cmdframework.provider.CommandNotification;
import com.ibm.websphere.management.async.client.AsyncCommandHandlerIF;

public class AsyncCmdTaskHandler implements AsyncCommandHandlerIF {

    public void handleNotification(CommandNotification notification) {
        // Add your own code here to handle the received notification
        System.out.println("\nEXAMPLE: notification received: " +
            notification);
    }
}

```

What to do next

You can complete other tasks associated with assets, such as deleting, editing, and exporting assets.

Viewing an asset using programming

You can view the asset information so that you can complete other tasks associated with the asset, such as editing or exporting an asset. An asset represents at least one binary file that implements business logic.

Before you begin

This task assumes a basic familiarity with command framework programming. Read about command framework programming in the application programming interface documentation.

Before you can view an asset of a business-level application, you must have imported an asset.

About this task

You can view an asset using programming, the administrative console, or the wsadmin tool. This topic describes how to view an asset using programming.

You must provide the assetID parameter to specify the asset you are viewing. You can view configuration information of an asset, such as the destination location and relationships with other assets.

Perform the following tasks to view an asset of a business-level application using programming.

Procedure

1. Connect to the application server.
The command framework allows the administrative command to be created and run with or without being connected to the application server. This step is optional if the application server is not running.
2. Create the command manager.
The command manager provides the functionality to create a new administrative command or query existing administrative commands.
3. Optionally create the asynchronous command handler for listening to command notifications.
Business-level application commands are implemented as asynchronous commands. To monitor the progress of the running command, you have to create an asynchronous command handler to receive notifications that the command generates.
4. Create the asynchronous command client.
An asynchronous command client provides a higher level interface to work with an asynchronous command. If you created an asynchronous command handler in the previous step, the handler is passed to the asynchronous command client. The asynchronous command client forwards the command notification to the handler and helps to control running of the command.
5. Use the command manager that you created in a previous step to create and set up the command to view an asset.
The command name is `viewAsset`. Use the required `assetID` parameter to specify the asset that you are viewing.
6. Call the `processCommandParameters` method in the asynchronous command client to process the command parameters.
The command framework asynchronous command model requires this call.
7. Call the asynchronous command client to run the command and view an asset.
You might have created an asynchronous command handler to implement the `AsyncCommandHandlerIF` interface class in a previous step. If you did, the asynchronous command client listens to command notifications and forwards the notifications to the handler. The handler performs any necessary actions while waiting for the command to complete.
8. Check the command result when the command completes.
When the command finishes running, control is returned to the caller. You can then check the result by calling the `command.getCommandResult` method.

Results

After you successfully run the code, you can view the configuration information of an asset.

Example

The following example shows how to view an asset based on the previous steps. Some statements are split on multiple lines for printing purposes.

```
package com.ibm.ws.management.application.task;

import java.util.Properties;

import com.ibm.websphere.management.AdminClient;
import com.ibm.websphere.management.AdminClientFactory;
import com.ibm.websphere.management.Session;
import com.ibm.websphere.management.cmdframework.AdminCommand;
import com.ibm.websphere.management.cmdframework.CommandMgr;
import com.ibm.websphere.management.cmdframework.CommandResult;
import com.ibm.websphere.management.cmdframework.CommandStep;
import com.ibm.websphere.management.cmdframework.TaskCommand;
import com.ibm.websphere.management.async.client.AsyncCommandClient;
```



```

public class ViewAsset {

    public static void main(String [] args) {

        try {
            // Connect to the application server.
            // This step is optional if you use the local
            // command manager. Comment out the lines to and including
            // CommandMgr cmdMgr = CommandMgr.getClientCommandMgr(
            // soapClient);
            // to get the soapClient soap client if you use the local
            // command manager.
            String host = "localhost";
            String port = "8880"; // Change to your port number if it is
                // not 8880.

            Properties config = new Properties();
            config.put(AdminClient.CONNECTOR_HOST, host);
            config.put(AdminClient.CONNECTOR_PORT, port);
            config.put(AdminClient.CONNECTOR_TYPE,
                AdminClient.CONNECTOR_TYPE_SOAP);
            System.out.println("Config: " + config);
            AdminClient soapClient =
                AdminClientFactory.createAdminClient(config);

            // Create the command manager.
            CommandMgr cmdMgr = CommandMgr.getClientCommandMgr(soapClient);

            // Comment out the previous lines to create a client command
            // manager if you are using a local command manager.
            // Uncomment the following line to create a local command
            // manager:
            //
            // CommandMgr cmdMgr = CommandMgr.getCommandMgr();
            System.out.println("\nCreated command manager");

            // Optionally create an asynchronous command handler.
            // Comment out the following line if no further handling
            // of command notification is required:
            AsyncCmdTaskHandler listener = new AsyncCmdTaskHandler();

            // Create an asynchronous command client.

            // Set up the session.
            String id = Long.toHexString(System.currentTimeMillis());
            String user = "content" + id;
            Session session = new Session(user, true);

            // If no command handler is used, replace the following listener with
            // null for the AsyncCommandClient object:
            AsyncCommandClient asyncCmdClientHelper = new
            AsyncCommandClient(session, listener);
            System.out.println("\nCreated async command client");

            // Create the command to view the asset.
            String cmdName = "viewAsset";
            AdminCommand cmd = cmdMgr.createCommand(cmdName);
            cmd.setConfigSession(session); // View a certain composition
                // unit of a business-level application
                // using the session created.
            System.out.println("\nCreated " + cmdName);

            // (required) Set the assetID parameter to the asset.
            // Examples of valid formats for the assetID parameter:
            // - aName
            // - assetname=aName
            // - WebSphere:assetname=aName

```

```

// This parameter accepts an incomplete ID as long as the
// incomplete ID can resolve to a unique asset.
String assetID = "asset1.zip";
cmd.setParameter("assetID", assetID);

System.out.println("\nSet assetID parameter to "
    + cmd.getParameter("assetID"));

// Call the asynchronous client helper to process parameters
try {
    asyncCmdClientHelper.processCommandParameters(cmd);
    System.out.println("\nCompleted process command " +
        "parameters");
} catch (Throwable th) {
    System.out.println("Failed from " +
        "asyncCmdClientHelper.processCommandParameters(cmd).");
    th.printStackTrace();
    System.exit(-1);
}

// Call the asynchronous command client to run the command.
asyncCmdClientHelper.execute(cmd);
System.out.println("\nCompleted running of command");

// Check the command result.
CommandResult result = cmd.getCommandResult();
if (result != null) {
    if (result.isSuccessful()) {
        System.out.println("\nCommand ran successfully "
            + "with result\n" + result.getResult());
    }
    else {
        System.out.println("\nCommand ran with " +
            "Exception");
        result.getException().printStackTrace();
    }
}
} catch (Exception e) {
    e.printStackTrace();
}
}

package com.ibm.ws.management.application.task;

import com.ibm.websphere.management.cmdframework.provider.CommandNotification;
import com.ibm.websphere.management.async.client.AsyncCommandHandlerIF;

public class AsyncCmdTaskHandler implements AsyncCommandHandlerIF {
    public void handleNotification(CommandNotification notification) {
        // Add your own code here to handle the received notification
        System.out.println("\nEXAMPLE: notification received: " +
            notification);
    }
}

```

What to do next

You can use the asset information that you viewed to perform other tasks. For instance, you might edit the asset to make improvements to the asset. You might export the asset and then import it into another configuration repository. You can then add the asset as a composition unit to a business-level application.

Editing an asset using programming

You can edit the information of an asset such as its destination location, its relationship with other assets, and so on. An asset represents at least one binary file that implements business logic.

Before you begin

This task assumes a basic familiarity with command framework programming. Read about command framework programming in the application programming interfaces documentation.

Before you can edit an asset, you must have imported an asset.

You can edit an asset of a business-level application using programming, the administrative console, or the wsadmin tool.

About this task

You can edit an asset of a business-level application using programming, the administrative console, or the wsadmin tool. This topic describes how to edit an asset of a business-level application using programming.

You must provide the `assetID` parameter to specify the asset that you are editing.

Perform the following tasks to edit an asset of a business-level application using programming.

Procedure

1. Connect to the application server.
The command framework allows the administrative command to be created and run with or without being connected to the application server. This step is optional if the application server is not running.
2. Create the command manager.
The command manager provides the functionality to create a new administrative command or query existing administrative commands.
3. Optionally create the asynchronous command handler for listening to command notifications.
Business-level application commands are implemented as asynchronous commands. To monitor the progress of the running command, you have to create an asynchronous command handler to receive notifications that the command generates.
4. Create the asynchronous command client.
An asynchronous command client provides a higher level interface to work with an asynchronous command. If you created an asynchronous command handler in the previous step, the handler is passed to the asynchronous command client. The asynchronous command client forwards the command notification to the handler and helps to control running of the command.
5. Use the command manager that you created in a previous step to create and set up the command that edits an asset.
The command name is `editAsset`. The `assetID` parameter is a required parameter to specify the asset that you are editing.
6. Call the asynchronous command client to process the command parameters.
7. Set up the command step parameters.
The `AssetOptions` step contains data about the asset such as its description, file permission, and relationship with other assets. You can edit various parameters in the `AssetOptions` step.
8. Call the asynchronous command client to run the command that edits an asset of a business-level application.

You might have created an asynchronous command handler to implement the `AsyncCommandHandlerIF` interface class in a previous step. If you did, the asynchronous command client listens to command notifications and forwards the notifications to the handler. The handler performs any necessary actions while waiting for the command to complete.

9. Check the command result when the command completes.

When the command finishes running, control is returned to the caller. You can then check the result by calling the `command.getCommandResult` method.

Results

After you successfully run the code, the asset of a business-level application is edited.

Example

The following example shows how to edit an asset of a business-level application based on the previous steps. Some statements are split on multiple lines for printing purposes.

```
package com.ibm.ws.management.application.task;

import java.util.Properties;

import com.ibm.websphere.management.AdminClient;
import com.ibm.websphere.management.AdminClientFactory;
import com.ibm.websphere.management.Session;
import com.ibm.websphere.management.cmdframework.AdminCommand;
import com.ibm.websphere.management.cmdframework.CommandMgr;
import com.ibm.websphere.management.cmdframework.CommandResult;
import com.ibm.websphere.management.cmdframework.CommandStep;
import com.ibm.websphere.management.cmdframework.TaskCommand;
import com.ibm.websphere.management.async.client.AsyncCommandClient;

public class EditAsset {

    public static void main(String [] args) {

        try {

            // Connect to the application server.
            // This step is optional if you use the local
            // command manager.

            // Comment out the lines to and including get the
            // soapClient soap client if you use the local command manager.
            // Comment out the lines to and including
            // CommandMgr cmdMgr = CommandMgr.getClientCommandMgr(
            // soapClient);
            // to get the soapClient soap client if you use the
            // local command manager.
            String host = "localhost";
            String port = "8880"; // Change to your port number if it is
                                // not 8880.

            Properties config = new Properties();
            config.put(AdminClient.CONNECTOR_HOST, host);
            config.put(AdminClient.CONNECTOR_PORT, port);
            config.put(AdminClient.CONNECTOR_TYPE,
                AdminClient.CONNECTOR_TYPE_SOAP);
            System.out.println("Config: " + config);
            AdminClient soapClient =
                AdminClientFactory.createAdminClient(config);

            // Create the command manager:
            CommandMgr cmdMgr = CommandMgr.getClientCommandMgr(soapClient);
```

```

// Comment out the previous lines to create a client command
// manager if you are using a local command manager.
// Uncomment the following line to create a local command
// manager.
// CommandMgr cmdMgr = CommandMgr.getCommandMgr();
System.out.println("\nCreated command manager");

// Optionally create async command handler.
// Comment out the following line if no further handling
// of command notification is required:
AsyncCmdTaskHandler listener = new AsyncCmdTaskHandler();

// Create an asynchronous command client.

// Set up the session.
String id = Long.toHexString(System.currentTimeMillis());
String user = "content" + id;
Session session = new Session(user, true);

// If no command handler is used, replace the listener with
// null for the following AsyncCommandClient object.
AsyncCommandClient asyncCmdClientHelper = new
AsyncCommandClient(session, listener);
System.out.println("\nCreated async command client");

// Create the command that edits the asset.
String cmdName = "editAsset";
AdminCommand cmd = cmdMgr.createCommand(cmdName);
cmd.setConfigSession(session); // Edit an imported asset
                                // using the session created.
System.out.println("\nCreated " + cmdName);

// Set the assetID parameter
// Examples of valid formats for the assetID parameter are:
// - aName
// - assetname=aName
// - WebSphere:assetname=aName
// This parameter accepts an incomplete ID as long as
// the incomplete ID can resolve to a unique asset.
String assetID = "asset1.zip";
cmd.setParameter("assetID", assetID);

System.out.println("\nSet assetID parameter to "
+ cmd.getParameter("assetID"));

// Call the asynchronous client helper to process parameters.
try {
    asyncCmdClientHelper.processCommandParameters(cmd);
    System.out.println("\nCompleted process command " +
"parameters");
} catch (Throwable th) {
    System.out.println("Failed from " +
"asyncCmdClientHelper.processCommandParameters(cmd).");
    th.printStackTrace();
    System.exit(-1);
}

// Set up the step parameters for the AssetOptions step.
String stepName = "AssetOptions";
CommandStep step = ((TaskCommand) cmd).gotoStep(stepName);

// Asset description:
String description = "asset for testing";

// destination of deployed asset
String destinationUrl = "/myInstalledAssets/asset1.zip";

```

```

// Asset type aspect:
String typeAspect = "spec=sharedlib";

// Asset validation:
String validate = "yes";

// Permission of files:
String filePermission = ".*\\.dll=755";

// Asset relationship:
String relationship = "";

for (int i = 0; i < step.getNumberOfRows(); i++) {
    // The following lines set the description and typeAspect
    // step parameters. There are other step parameters
    // in the AssetOptions step in the following comments. Change your set
    // of step parameters as required by your scenario.

    // For example, set description
    step.setParameter("description", description, i);
    System.out.println("\nSet description parameter to " +
        step.getParameter("description", i));

    // For example, set the typeAspect parameter.
    // Format of a typeAspect is:
    // WebSphere:spec=xxx,version=n.n+
    // WebSphere:spec=xxx,version=n.n
    step.setParameter("typeAspect", typeAspect, i);
    System.out.println("\nGet typeAspect: " +
        step.getParameter("typeAspect", i));

    // For example, set the destination parameter.
    step.setParameter("destination", destination, i);
    System.out.println("\nSet destination parameter to " +
        step.getParameter("destination", i));

    // For example, set the validate parameter.
    step.setParameter("validate", validate, i);
    System.out.println("\nSet validate parameter to " +
        step.getParameter("validate", i));

    // For example, set the filePermission parameter.
    step.setParameter("filePermission", filePermission, i);
    System.out.println("\nSet filePermission parameter to " +
        step.getParameter("filePermission", i));

    // For example, set relationship.
    step.setParameter("relationship", relationship, i);
    System.out.println("\nSet relationship paramter to " +
        step.getParameter("relationship", i));
}

// Run the command to edit the asset.
asyncCmdClientHelper.execute(cmd);
System.out.println("\nCompleted running of command");

// Check the command result.
CommandResult result = cmd.getCommandResult();
if (result != null) {
    if (result.isSuccessful()) {
        System.out.println("\nCommand ran successfully "
            + "with result\n" + result.getResult());
    }
    else {
        System.out.println("\nCommand ran with " +
            "Exception");
        result.getException().printStackTrace();
    }
}

```

```

        }
    } catch (Exception e) {
        e.printStackTrace();
    }
}

package com.ibm.ws.management.application.task;

import com.ibm.websphere.management.cmdframework.provider.CommandNotification;
import com.ibm.websphere.management.async.client.AsyncCommandHandlerIF;

public class AsyncCmdTaskHandler implements AsyncCommandHandlerIF {

    public void handleNotification(CommandNotification notification) {
        // Add your own code here to handle the received notification
        System.out.println("\nEXAMPLE: notification received: " +
            notification);
    }
}

```

What to do next

After you edit the asset, you can add the asset as a composition unit to a business-level application, or export the asset.

Deleting an asset using programming

You can delete an asset from a business-level application using programming if the asset is not functioning correctly, the asset is no longer needed, and so on. An asset represents at least one binary file that implements business logic.

Before you begin

This task assumes a basic familiarity with command framework programming. Read about command framework programming in the application programming interface documentation.

Before you can delete an asset, you must have imported the asset. All the composition units associated with the asset must be deleted using the `deleteCompUnit` command before you delete the asset. Otherwise, you have to force the deletion. If you do not force the deletion, the deletion fails. If any other composition units have a dependency on a composition unit being deleted with the `force` option, the deletion fails. After all dependencies on the composition unit are removed, the `force` option succeeds.

About this task

You can delete an asset using programming, the administrative console, or the `wsadmin` tool. Use this topic to delete an asset using programming.

You must specify the `assetID` parameter of the asset that you are deleting. You might delete an asset if it is not functioning correctly, it is no longer needed, and so on.

Perform the following tasks to delete an asset using programming.

Procedure

1. Connect to the application server.

The command framework allows the administrative command to be created and run with or without being connected to the application server. This step is optional if the application server is not running.

2. Create the command manager.
The command manager provides the functionality to create a new administrative command or query existing administrative commands.
3. Optionally create the asynchronous command handler for listening to command notifications.
Business-level application commands are implemented as asynchronous commands. To monitor the progress of the running command, you have to create an asynchronous command handler to receive notifications that the command generates.
4. Create the asynchronous command client.
An asynchronous command client provides a higher level interface to work with an asynchronous command. If you created an asynchronous command handler in the previous step, the handler is passed to the asynchronous command client. The asynchronous command client forwards the command notification to the handler and helps to control running of the command.
5. Use the command manager that you created in a previous step to create and set up the command that deletes an asset.
The command name is `deleteAsset`. The `assetID` parameter is a required parameter to specify the asset to delete. You can optionally specify the `delete` parameter to force deletion of an asset if composition units are still associated with the asset.
6. Call the `processCommandParameters` method in the asynchronous command client to process the command parameters.
The command framework asynchronous command model requires this call.
7. Call the asynchronous command client to run the command that deletes an asset.
You might have created an asynchronous command handler to implement the `AsyncCommandHandlerIF` interface class in a previous step. If you did, the asynchronous command client listens to command notifications and forwards the notifications to the handler. The handler performs any necessary actions while waiting for the command to complete.
8. Check the command result when the command completes.
When the command finishes running, control is returned to the caller. You can then check the result by calling the `command.getCommandResult` method.

Results

After you successfully run the code, the asset is deleted.

Example

The following example shows how to delete an asset from a business-level application based on the previous steps. Some statements are split on multiple lines for printing purposes.

```
package com.ibm.ws.management.application.task;

import java.util.Properties;

import com.ibm.websphere.management.AdminClient;
import com.ibm.websphere.management.AdminClientFactory;
import com.ibm.websphere.management.Session;
import com.ibm.websphere.management.cmdframework.AdminCommand;
import com.ibm.websphere.management.cmdframework.CommandMgr;
import com.ibm.websphere.management.cmdframework.CommandResult;
import com.ibm.websphere.management.cmdframework.CommandStep;
import com.ibm.websphere.management.cmdframework.TaskCommand;
import com.ibm.websphere.management.async.client.AsyncCommandClient;

public class DeleteAsset {

    public static void main(String [] args) {

        try {
```



```

// Connect to the application server.
// This step is optional if you use the local
// command manager. Comment out the lines to and including
// CommandMgr cmdMgr = CommandMgr.getClientCommandMgr(
// soapClient);
// to get the soapClient soap client if you use the local
// command manager.

String host = "localhost";
String port = "8880"; // Change to your port number if it is
                    // not 8880.

Properties config = new Properties();
config.put(AdminClient.CONNECTOR_HOST, host);
config.put(AdminClient.CONNECTOR_PORT, port);
config.put(AdminClient.CONNECTOR_TYPE,
           AdminClient.CONNECTOR_TYPE_SOAP);
System.out.println("Config: " + config);
AdminClient soapClient =
    AdminClientFactory.createAdminClient(config);

// Create the command manager.
CommandMgr cmdMgr = CommandMgr.getClientCommandMgr(soapClient);

// Comment out the previous lines to create a client command
// manager if you are using a local command manager.
// Uncomment the following line to create a local command
// manager.
//
// CommandMgr cmdMgr = CommandMgr.getCommandMgr();
System.out.println("\nCreated command manager");

// Optionally create an asynchronous command handler.
// Comment out the following line if no further handling
// of command notification is required.
AsyncCmdTaskHandler listener = new AsyncCmdTaskHandler();

// Create an asynchronous command client.

// Set up the session.
String id = Long.toHexString(System.currentTimeMillis());
String user = "content" + id;
Session session = new Session(user, true);

// If no command handler is used, replace listener with
// null for the following AsyncCommandClient object:
AsyncCommandClient asyncCmdClientHelper = new
AsyncCommandClient(session, listener);
System.out.println("\nCreated async command client");

// Create the command that deletes the asset.
String cmdName = "deleteAsset";
AdminCommand cmd = cmdMgr.createCommand(cmdName);
cmd.setConfigSession(session); // Delete the asset from the
// business-level application using the session created.
System.out.println("\nCreated " + cmdName);

// Set the assetID parameter to the asset that is to be
// deleted.
// Examples of valid formats for the assetID parameter are:
// - aName
// - assetname=aName
// - WebSphere:assetname=aName
// This parameter will accept an incomplete ID as long as
// the incomplete ID can resolve to a unique asset

```

```

// in the business-level application.
String assetID = "as1";
cmd.setParameter("assetID", assetID );

System.out.println("\nSet assetID parameter to "
    + cmd.getParameter("assetID"));
// Uncomment the following line of code to set the force parameter
// to force the deletion even if there are composition units
// associated with this asset.
//
// cmd.setParameter("force", "true");

// Call the asynchronous client helper to process parameters.
try {
    asyncCmdClientHelper.processCommandParameters(cmd);
    System.out.println("\nCompleted process command " +
        "parameters");
} catch (Throwable th) {
    System.out.println("Failed from " +
        "asyncCmdClientHelper.processCommandParameters(cmd).");
    th.printStackTrace();
    System.exit(-1);
}

// Call the asynchronous command client to run the command.
asyncCmdClientHelper.execute(cmd);
System.out.println("\nCompleted running the command");

// Check the command result.
CommandResult result = cmd.getCommandResult();
if (result != null) {
    if (result.isSuccessful()) {
        System.out.println("\nCommand ran successfully "
            + "with result\n" + result.getResult());
    }
    else {
        System.out.println("\nCommand ran with " +
            "Exception");
        result.getException().printStackTrace();
    }
}
} catch (Exception e) {
    e.printStackTrace();
}
}
}
}

```

```

package com.ibm.ws.management.application.task;

import com.ibm.websphere.management.cmdframework.provider.CommandNotification;
import com.ibm.websphere.management.async.client.AsyncCommandHandlerIF;

public class AsyncCmdTaskHandler implements AsyncCommandHandlerIF {

    public void handleNotification(CommandNotification notification) {
        // Add your own code here to handle the received notification
        System.out.println("\nEXAMPLE: notification received: " +
            notification);
    }
}

```

What to do next

You can complete other steps associated with assets in business-level applications, such as adding or deleting other assets, listing assets, exporting assets, and so on.

Exporting an asset using programming

You can export an asset from the current session so that you can back up the asset, import the asset to another session, and so on. An asset represents at least one binary file that implements business logic.

Before you begin

This task assumes a basic familiarity with command framework programming. Read about command framework programming in the application programming interfaces documentation.

This task assumes that you have already imported an asset.

About this task

You can export an asset using programming, the administrative console, or the wsadmin tool. This topic describes how to export an asset using programming.

You must provide an `assetID` parameter value and a file name parameter value to export an asset. The `assetID` parameter identifies the asset you want to export. An asset ID can take a number of forms. The list below shows various forms for an asset named `asset1.jar`.

- `asset1.jar`
- `assetname=asset1.jar`
- `WebSphere:assetname=asset1.jar`

The `filename` parameter specifies a file system file name and location for the exported asset. Specify a fully qualified file path for the file name parameter because the results with relative path names are unpredictable. If you specify a file name parameter of a file that already exists, the file is overwritten with the exported asset.

Perform the following tasks to export an asset from a business-level application using programming.

Procedure

1. Connect to the application server.

The command framework allows the administrative command to be created and run with or without being connected to the application server. This step is optional if the application server is not running.

2. Create the command manager.

The command manager provides the functionality to create a new administrative command or query existing administrative commands.

3. Optionally create the asynchronous command handler for listening to command notifications.

Business-level application commands are implemented as asynchronous commands. To monitor the progress of the running command, you have to create an asynchronous command handler to receive notifications that the command generates.

4. Create the asynchronous command client.

An asynchronous command client provides a higher level interface to work with an asynchronous command. If you created an asynchronous command handler in the previous step, the handler is passed to the asynchronous command client. The asynchronous command client forwards the command notification to the handler and helps to control running of the command.

5. Use the command manager that you created in a previous step to create and set up the command that exports an asset.

The command name is `exportAsset`. The `assetID` and `filename` parameters are required parameters to specify the asset to export and the file name and directory where the asset is exported.

6. Call the `processCommandParameters` method in the asynchronous command client to process the command parameters.

The command framework asynchronous command model requires this call.

7. Call the asynchronous command client to run the command that exports an asset.

You might have created an asynchronous command handler to implement the `AsyncCommandHandlerIF` interface class in a previous step. If you did, the asynchronous command client listens to command notifications and forwards the notifications to the handler. The handler performs any necessary actions while waiting for the command to complete.

8. Check the command result when the command completes.

When the command finishes running, control is returned to the caller. You can then check the result by calling the `command.getCommandResult` method.

Results

After you successfully run the code, the asset is exported.

Example

The following example shows how to export an asset from a business-level application based on the previous steps.

Some statements are split on multiple lines for printing purposes.

```
package com.ibm.ws.management.application.task;

import java.util.Properties;

import com.ibm.websphere.management.AdminClient;
import com.ibm.websphere.management.AdminClientFactory;
import com.ibm.websphere.management.Session;
import com.ibm.websphere.management.cmdframework.AdminCommand;
import com.ibm.websphere.management.cmdframework.CommandMgr;
import com.ibm.websphere.management.cmdframework.CommandResult;
import com.ibm.websphere.management.cmdframework.TaskCommand;
import com.ibm.websphere.management.async.client.AsyncCommandClient;

public class ExportAsset {

    public static void main(String [] args) {

        try {

            // Connect to the application server.
            // This step is optional if you use the local
            // command manager. Comment out the lines to and including
            // CommandMgr cmdMgr = CommandMgr.getClientCommandMgr(
            // soapClient);
            // to get the soapClient soap client if you use the local
            // command manager.

            String host = "localhost";
            String port = "8880"; // Change to your port number if it is
            // not 8880.

            Properties config = new Properties();
            config.put(AdminClient.CONNECTOR_HOST, host);
            config.put(AdminClient.CONNECTOR_PORT, port);
            config.put(AdminClient.CONNECTOR_TYPE,
                AdminClient.CONNECTOR_TYPE_SOAP);
            System.out.println("Config: " + config);
            AdminClient soapClient =
                AdminClientFactory.createAdminClient(config);
```

```

// Create the command manager.
CommandMgr cmdMgr = CommandMgr.getClientCommandMgr(soapClient);

// Comment out the previous lines to create a client command
// manager if you are using a local command manager.
// Uncomment the following line to create a local command
// manager:
//
// CommandMgr cmdMgr = CommandMgr.getCommandMgr();
System.out.println("\nCreated command manager");

// Optionally create an asynchronous command handler.
// Comment out the following line if no further handling
// of command notification is required:
AsyncCmdTaskHandler listener = new AsyncCmdTaskHandler();

// Create an asynchronous command client.

// Set up the session.
String id = Long.toHexString(System.currentTimeMillis());
String user = "content" + id;
Session session = new Session(user, true);

// If no command handler is used, replace the following listener with
// null for the AsyncCommandClient object.
AsyncCommandClient asyncCmdClientHelper = new
AsyncCommandClient(session, listener);
System.out.println("\nCreated async command client");

// Create command that exports the asset.
String cmdName = "exportAsset";
AdminCommand cmd = cmdMgr.createCommand(cmdName);
cmd.setConfigSession(session); // Export as asset
// using the session created.
System.out.println("\nCreated " + cmdName);

// (required) Set the assetID parameter to the composition
// unit that you are exporting.
// Examples of valid formats for the assetID parameter are:
// - aName
// - assetname=aName
// - WebSphere:assetname=aName
// This parameter accepts an incomplete ID as long as
// the incomplete ID can resolve to a unique asset
// within the business-level application.
String assetID = "test5.zip";
cmd.setParameter("assetID", assetID);

System.out.println("\nSet assetID parameter to "
+ cmd.getParameter("assetID"));

// Set the file name for the asset to be exported. Use a
// fully qualified path name. An existing file with the specified
// name will be overwritten.
DownloadFile filename = new DownloadFile("/assets/asset1.zip");

cmd.setParameter("filename", filename);

System.out.println("\nSet filename parameter to "
+ cmd.getParameter("filename"));

// Call the asynchronous client helper to process parameters.
try {
    asyncCmdClientHelper.processCommandParameters(cmd);
}

```

```

        System.out.println("\nCompleted process command " +
            "parameters");
    } catch (Throwable th) {
        System.out.println("Failed from " +
            "asyncCmdClientHelper.processCommandParameters(cmd).");
        th.printStackTrace();
        System.exit(-1);
    }

    // Call the asynchronous command client to run the command.
    asyncCmdClientHelper.execute(cmd);
    System.out.println("\nCompleted running of the command");

    // Check the command result.
    CommandResult result = cmd.getCommandResult();
    if (result != null) {
        if (result.isSuccessful()) {
            System.out.println("\nCommand ran successfully "
                + "with result\n" + result.getResult());
        }
        else {
            System.out.println("\nCommand ran with " +
                "Exception");
            result.getException().printStackTrace();
        }
    }
} catch (Exception e) {
    e.printStackTrace();
}
}
}

package com.ibm.ws.management.application.task;

import com.ibm.websphere.management.cmdframework.provider.CommandNotification;
import com.ibm.websphere.management.async.client.AsyncCommandHandlerIF;

public class AsyncCmdTaskHandler implements AsyncCommandHandlerIF {

    public void handleNotification(CommandNotification notification) {
        // Add your own code here to handle the received notification
        System.out.println("\nEXAMPLE: notification received: " +
            notification);
    }
}
}

```

What to do next

You can import the asset to another session. You can complete other tasks associated with assets, such as listing assets, and editing assets.

Starting a business-level application using programming

You can start a business-level application, which starts each composition unit in that business-level application. Each composition unit is started on the respective targets on which the business-level application is deployed.

Before you begin

Before you can start a business-level application, you must have created an empty business-level application, imported an asset, and added a composition unit to the business-level application.

You can start a business-level application using programming, the administrative console, or the wsadmin tool.

This task assumes a basic familiarity with command framework programming. Read about command framework programming in the application programming interfaces documentation.

About this task

You must specify the blaID parameter of the business-level application that you start.

Perform the following steps to start a business-level application using programming.

Procedure

1. Connect to the application server.
The command framework allows the administrative command to be created and run with or without being connected to the application server. This step is optional if the application server is not running.
2. Create the command manager.
The command manager provides the functionality to create a new administrative command or query existing administrative commands.
3. Optionally create the asynchronous command handler for listening to command notifications.
Business-level application commands are implemented as asynchronous commands. To monitor the progress of the running command, you have to create an asynchronous command handler to receive notifications that the command generates.
4. Create the asynchronous command client.
An asynchronous command client provides a higher level interface to work with an asynchronous command. If you created an asynchronous command handler in the previous step, the handler is passed to the asynchronous command client. The asynchronous command client forwards the command notification to the handler and helps to control running of the command.
5. Use the command manager that you created in a previous step to set up the command that starts a business-level application.
The command name is startBLA. The blaID parameter is a required parameter to specify the business-level application to start.
6. Call the processCommandParameters method in the asynchronous command client to process the command parameters.
The command framework asynchronous command model requires this call.
7. Call the asynchronous command client to run the command that starts a business-level application.
You might have created an asynchronous command handler to implement the AsyncCommandHandlerIF interface class in a previous step. If you did, the asynchronous command client listens to command notifications and forwards the notifications to the handler. The handler performs any necessary actions while waiting for the command to complete.
8. Check the command result when the command completes.
When the command finishes running, control is returned to the caller. You can then check the result by calling the command.getCommandResult method.

Results

After you successfully run the code, the business-level application is started.

Example

The following example shows how to start a business-level application based on the previous steps. Some statements are split on multiple lines for printing purposes.

```

package com.ibm.ws.management.application.task;

import java.util.Properties;

import com.ibm.websphere.management.AdminClient;
import com.ibm.websphere.management.AdminClientFactory;
import com.ibm.websphere.management.Session;
import com.ibm.websphere.management.cmdframework.AdminCommand;
import com.ibm.websphere.management.cmdframework.CommandMgr;
import com.ibm.websphere.management.cmdframework.CommandResult;
import com.ibm.websphere.management.cmdframework.CommandStep;
import com.ibm.websphere.management.cmdframework.TaskCommand;
import com.ibm.websphere.management.async.client.AsyncCommandClient;

public class startBLA {

    public static void main(String [] args) {

        try {

            // Connect to the application server.
            // This step is optional if you use the local command
            // manager. Comment out the lines to and including
            // CommandMgr cmdMgr = CommandMgr.getClientCommandMgr(
            // soapClient);
            // to get the soapClient soap client if
            // you use the local command manager.

            String host = "localhost";
            String port = "8880"; // Change to your port number if it is
            // not 8880.

            Properties config = new Properties();
            config.put(AdminClient.CONNECTOR_HOST, host);
            config.put(AdminClient.CONNECTOR_PORT, port);
            config.put(AdminClient.CONNECTOR_TYPE,
                AdminClient.CONNECTOR_TYPE_SOAP);
            System.out.println("Config: " + config);
            AdminClient soapClient =
                AdminClientFactory.createAdminClient(config);

            // Create the command manager
            CommandMgr cmdMgr = CommandMgr.getClientCommandMgr(soapClient);

            // Comment out the previous lines to create a client command
            // manager if you are using a local command manager.
            // Uncomment the following line to create a local command
            // manager:
            //
            // CommandMgr cmdMgr = CommandMgr.getCommandMgr();
            System.out.println("\nCreated command manager");

            // Optionally create an asynchronous command handler
            // for listening to command notifications.
            // Comment out the following line if no further handling
            // of command notification is required:
            AsyncCmdTaskHandler listener = new AsyncCmdTaskHandler();

            // Create an asynchronous command client.

            // Set up the session.
            String id = Long.toHexString(System.currentTimeMillis());
            String user = "content" + id;
            Session session = new Session(user, true);

            // If no command handler is used, replace the listener with
            // null for the AsyncCommandClient object that follows.

```



```

AsyncCommandClient asyncCmdClientHelper = new
AsyncCommandClient(session, listener);
System.out.println("\nCreated async command client");

// Create the command that starts the business-level application.
String cmdName = "startBLA";
AdminCommand cmd = cmdMgr.createCommand(cmdName);
cmd.setConfigSession(session); // Start a business-level
//application using the session created.
System.out.println("\nCreated " + cmdName);

// (required) Set the blaID parameter.
// Examples of valid formats for the blaID parameter are:
// - bName
// - blaname=bName
// - WebSphere:blaname=bName
// This parameter
// accepts an incomplete ID as long as the incomplete
// ID can resolve to a unique business-level application.
String blaID = "bla1";
cmd.setParameter("blaID", blaID);

System.out.println("\nSet blaID parameter to "
+ cmd.getParameter("blaID"));

// Call asynchronous client helper to process parameters.
try {
    asyncCmdClientHelper.processCommandParameters(cmd);
    System.out.println("\nCompleted process command " +
        "parameters");
} catch (Throwable th) {
    System.out.println("Failed from " +
        "asyncCmdClientHelper.processCommandParameters(cmd).");
    th.printStackTrace();
    System.exit(-1);
}

// Call the asynchronous command client to run the command command.
asyncCmdClientHelper.execute(cmd);
System.out.println("\nCompleted running the command");

// Check the command result.
CommandResult result = cmd.getCommandResult();
if (result != null) {
    if (result.isSuccessful()) {
        System.out.println("\nCommand ran successfully "
+ "with result\n" + result.getResult());
    }
    else {
        System.out.println("\nCommand ran with " +
            "Exception");
        result.getException().printStackTrace();
    }
}
} catch (Exception e) {
    e.printStackTrace();
}
}

package com.ibm.ws.management.application.task;

import com.ibm.websphere.management.cmdframework.provider.CommandNotification;
import com.ibm.websphere.management.async.client.AsyncCommandHandlerIF;

public class AsyncCmdTaskHandler implements AsyncCommandHandlerIF {

```

```
public void handleNotification(CommandNotification notification) {  
    // Add your own code here to handle the received notification  
    System.out.println("\nEXAMPLE: notification received: " +  
        notification);  
}  
}
```

What to do next

Your users can access the business-level application that you started.

Stopping a business-level application using programming

You can stop a business-level application, which stops each composition unit in that business-level application. Each composition unit is stopped on the respective targets on which the business-level application is deployed.

Before you begin

Before you can stop a business-level application, you must have created an empty business-level application, imported an asset, added a composition unit to the business-level application, and started the business-level application.

About this task

You can stop a business-level application using programming, the administrative console, or the wsadmin tool. This topic describes how to stop a business-level application using programming.

Perform the following steps to stop a business-level application using programming.

Procedure

1. Connect to the application server.
The command framework allows the administrative command to be created and run with or without being connected to the application server. This step is optional if the application server is not running.
2. Create the command manager.
The command manager provides the functionality to create a new administrative command or query existing administrative commands.
3. Optionally create the asynchronous command handler for listening to command notifications.
Business-level application commands are implemented as asynchronous commands. To monitor the progress of the running command, you have to create an asynchronous command handler to receive notifications that the command generates.
4. Create the asynchronous command client.
An asynchronous command client provides a higher level interface to work with an asynchronous command. If you created an asynchronous command handler in the previous step, the handler is passed to the asynchronous command client. The asynchronous command client forwards the command notification to the handler and helps to control running of the command.
5. Use the command manager that you created in a previous step Create and set up the command that stops a business-level application.
The command name is stopBLA. The blaID parameter is a required parameter to specify the business-level application to stop.
6. Call the processCommandParameters method in the asynchronous command client to process the command parameters.
The command framework asynchronous command model requires this call.

7. Call the asynchronous command client to run the command that stops a business-level application. You might have created an asynchronous command handler to implement the `AsyncCommandHandlerIF` interface class in a previous step. If you did, the asynchronous command client listens to command notifications and forwards the notifications to the handler. The handler performs any necessary actions while waiting for the command to complete.
8. Check the command result when the command completes. When the command finishes running, control is returned to the caller. You can then check the result by calling the `command.getCommandResult` method.

Results

After you successfully run the code, the business-level application is stopped.

Example

The following example shows how to stop a business-level application based on the previous steps. Some statements are split on multiple lines for printing purposes.

```
package com.ibm.ws.management.application.task;

import java.util.Properties;

import com.ibm.websphere.management.AdminClient;
import com.ibm.websphere.management.AdminClientFactory;
import com.ibm.websphere.management.Session;
import com.ibm.websphere.management.cmdframework.AdminCommand;
import com.ibm.websphere.management.cmdframework.CommandMgr;
import com.ibm.websphere.management.cmdframework.CommandResult;
import com.ibm.websphere.management.cmdframework.CommandStep;
import com.ibm.websphere.management.cmdframework.TaskCommand;
import com.ibm.websphere.management.async.client.AsyncCommandClient;

public class stopBLA {

    public static void main(String [] args) {

        try {

            // Connect to the application server.
            // This step is optional if you use the local command
            // manager. Comment out the lines to and including
            // CommandMgr cmdMgr = CommandMgr.getClientCommandMgr(
            // soapClient);
            // to get the soapClient soap client if
            // you use the local command manager.

            String host = "localhost";
            String port = "8880"; // Change to your port number if it is
                // not 8880.

            Properties config = new Properties();
            config.put(AdminClient.CONNECTOR_HOST, host);
            config.put(AdminClient.CONNECTOR_PORT, port);
            config.put(AdminClient.CONNECTOR_TYPE,
                AdminClient.CONNECTOR_TYPE_SOAP);
            System.out.println("Config: " + config);
            AdminClient soapClient =
                AdminClientFactory.createAdminClient(config);

            // Create the command manager
            CommandMgr cmdMgr = CommandMgr.getClientCommandMgr(soapClient);

            // Comment out the previous lines to create a client command
            // manager if you are using a local command manager.
```

```

// Uncomment the following line to create a local command
// manage:.
//
// CommandMgr cmdMgr = CommandMgr.getCommandMgr();
System.out.println("\nCreated command manager");

// Optionally create an asynchronous command handler
// for listening to command notifications.
// Comment out the following line if no further handling
// of command notification is required.
AsyncCmdTaskHandler listener = new AsyncCmdTaskHandler();

// Create an asynchronous command client.

// Set up the session.
String id = Long.toHexString(System.currentTimeMillis());
String user = "content" + id;
Session session = new Session(user, true);

// If no command handler is used, replace the listener with
// null for the AsyncCommandClient object that follows.
AsyncCommandClient asyncCmdClientHelper = new
AsyncCommandClient(session, listener);
System.out.println("\nCreated async command client");

// Create the command that stops the business-level application.
String cmdName = "stopBLA";
AdminCommand cmd = cmdMgr.createCommand(cmdName);
cmd.setConfigSession(session); // Stop a business-level
//application that is using the session created.
System.out.println("\nCreated " + cmdName);

// (required) Set the blaID parameter.
// Examples of valid formats for the blaID parameter are:
// - bName
// - blaname=bName
// - WebSphere:blaname=bName
// This parameter
// accepts an incomplete ID as long as the incomplete
// ID can resolve to a unique business-level application.
String blaID = "bla1";
cmd.setParameter("blaID", blaID);

System.out.println("\nSet blaID parameter to "
+ cmd.getParameter("blaID"));

// Call asynchronous client helper to process parameters.
try {
    asyncCmdClientHelper.processCommandParameters(cmd);
    System.out.println("\nCompleted process command " +
        "parameters");
} catch (Throwable th) {
    System.out.println("Failed from " +
        "asyncCmdClientHelper.processCommandParameters(cmd).");
    th.printStackTrace();
    System.exit(-1);
}

// Call the asynchronous command client to run the command command.
asyncCmdClientHelper.execute(cmd);
System.out.println("\nCompleted running of command");

// Check the command result.
CommandResult result = cmd.getCommandResult();
if (result != null) {
    if (result.isSuccessful()) {
        System.out.println("\nCommand ran successfully "

```

```

        + "with result\n" + result.getResult());
    }
    else {
        System.out.println("\nCommand ran with " +
            "Exception");
        result.getException().printStackTrace();
    }
} catch (Exception e) {
    e.printStackTrace();
}
}
}

package com.ibm.ws.management.application.task;

import com.ibm.websphere.management.cmdframework.provider.CommandNotification;
import com.ibm.websphere.management.async.client.AsyncCommandHandlerIF;

public class AsyncCmdTaskHandler implements AsyncCommandHandlerIF {

    public void handleNotification(CommandNotification notification) {
        // Add your own code here to handle the received notification
        System.out.println("\nEXAMPLE: notification received: " +
            notification);
    }
}

```

What to do next

Complete administrative tasks on the business-level application, such as editing an asset or a composition unit that is contained in the business-level application.

Checking the status of a business-level application using programming

You can check the status of an entire business-level application. You can also limit the status to a particular composition unit of a business-level application, a specific deployment target, or check the status of the composition unit and the deployment target at the same time.

Before you begin

This task assumes a basic familiarity with command framework programming. Read about command framework programming in the application programming interface documentation.

Before you can check the status of a business-level application or a composition unit, you must have created the business-level application.

You can check the status of a business-level application using programming, the administrative console, or the wsadmin tool.

About this task

You must provide the blaID parameter to specify the business-level application that you are viewing.

Perform the following tasks to view a business-level application using programming.

Procedure

1. Connect to the application server.

The command framework allows the administrative command to be created and run with or without being connected to the application server. This step is optional if the application server is not running.

2. Create the command manager.

The command manager provides the functionality to create a new administrative command or query existing administrative commands.

3. Optionally create the asynchronous command handler for listening to command notifications.

Business-level application commands are implemented as asynchronous commands. To monitor the progress of the running command, you have to create an asynchronous command handler to receive notifications that the command generates.

4. Create the asynchronous command client.

An asynchronous command client provides a higher level interface to work with an asynchronous command. If you created an asynchronous command handler in the previous step, the handler is passed to the asynchronous command client. The asynchronous command client forwards the command notification to the handler and helps to control running of the command.

5. Create and set up the `getBLAStatus` command to check the status of a business-level application.

- a. Set the `blaID` parameter for the business-level application whose status you want to check.
- b. Optionally set the `culD` parameter if you want to narrow the scope of the query to a single composition unit.
- c. Optionally set the `targetID` if you want to narrow the scope of the query to a single target server process or cluster.

6. Call the `processCommandParameters` method in the asynchronous command client to process the command parameters.

The command framework asynchronous command model requires this call.

7. Call the asynchronous command client to run the command to check the status of the business-level application.

You could have created an asynchronous command handler to implement the `AsyncCommandHandlerIF` interface class in a previous step. If you did, the asynchronous command client listens to command notifications and forwards the notifications to the handler. The handler performs any necessary actions while waiting for the command to complete.

8. Check the command result when the command completes.

When the command finishes running, control is returned to the caller. You can then check the result by calling the `command.getCommandResult` method.

Results

After you successfully run the code, you can check the status of an entire business-level application, if you chose not to limit the status. If you chose options to limit the status, you could check the status to a particular composition unit of a business-level application, a specific deployment target, or check the status of the composition unit and the deployment target at the same time.

The smallest unit of status data that the system maintains is for a single composition unit in a single server or cluster member process. Business-level application status can be based on one or more composition units, each having one or more targets, with targets potentially consisting of clusters with multiple member processes. Therefore, the single status value returned from the `getBLAStatus` command is a compilation of individual status data for all composition units on all target process within the scope of the status query. The following table describes how individual status data is compiled into a single status value. The term *composition unit instance* used in the table refers to a composition unit on a single server or single cluster member process.

Table 71. Business-level application status descriptions. Read the descriptions to learn about application status.

Status	Description
ExecutionState.STARTED	All composition unit instances within the scope of the query have been started.
ExecutionState.STOPPED	All composition unit instances within the scope of the query have not been started or have been stopped.
ExecutionState.PARTIAL_START	Some composition unit instances within the scope of the query have a status of ExecutionState.STARTED and some have a status of ExecutionState.STOPPED.
ExecutionState.UNKNOWN	Status data for at least one composition instance within the scope of the query cannot be obtained for some reason.

Example

The following example shows how to check the status of a business-level application based on the previous steps. Some statements are split on multiple lines for printing purposes.

```
package com.ibm.ws.management.application.task;

import java.util.Properties;

import com.ibm.websphere.management.AdminClient;
import com.ibm.websphere.management.AdminClientFactory;
import com.ibm.websphere.management.Session;
import com.ibm.websphere.management.cmdframework.AdminCommand;
import com.ibm.websphere.management.cmdframework.CommandMgr;
import com.ibm.websphere.management.cmdframework.CommandResult;
import com.ibm.websphere.management.cmdframework.CommandStep;
import com.ibm.websphere.management.cmdframework.TaskCommand;
import com.ibm.websphere.management.async.client.AsyncCommandClient;

public class EditBLA {

    public static void main(String[] args) {

        try {

            // Connect to the application server.
            // This step is optional if you use the local command manager.
            // Comment out the following lines to get the soapClient soap client if
            // you are going to use the local command manager. You would
            // comment out the lines to and including
            // CommandMgr cmdMgr =
            // CommandMgr.getClientCommandMgr(soapClient);

            String host = "localhost"; // Change to your host if it is not localhost.
            String port = "8880"; // Change to your port number if it is not 8880.

            Properties config = new Properties();
            config.put(AdminClient.CONNECTOR_HOST, host);
            config.put(AdminClient.CONNECTOR_PORT, port);
            config.put(AdminClient.CONNECTOR_TYPE,
                AdminClient.CONNECTOR_TYPE_SOAP);
            System.out.println("Config: " + config);
            AdminClient soapClient =
                AdminClientFactory.createAdminClient(config);

            // Create command manager.
            CommandMgr cmdMgr = CommandMgr.getClientCommandMgr(soapClient);

            // Comment out the previous lines to create a client command
            // manager if you are using a local command manager.
            // Uncomment the following line to create a local command
            // manager.
            //
            // CommandMgr cmdMgr = CommandMgr.getCommandMgr();

            System.out.println("\nCreated command manager");

            // Optionally create an asynchronous command handler.
            // Comment out the following line if no further handling
            // of command notification is required.
            AsyncCmdTaskHandler listener = new AsyncCmdTaskHandler();

            // Create an asynchronous command client.

            // Set up the session.
            // This example creates a new session. You can replace the
            // following code to use an existing session that has been
            // created.
```

```

String id = Long.toHexString(System.currentTimeMillis());
String user = "content" + id;
Session session = new Session(user, true);

// If no command handler is used, replace the listener with
// null for the following AsyncCommandClient object.
AsyncCommandClient asyncCmdClientHelper = new
    AsyncCommandClient(session, listener);
System.out.println("\nCreated async command client");

// Create the command.
String cmdName = "getBLAstatus";
AdminCommand cmd = cmdMgr.createCommand(cmdName);
cmd.setConfigSession(session); // Check the status
                                // using the session
                                // created
System.out.println("\nCreated " + cmdName);

// Set the required blaID parameter
// Examples of valid formats for the blaID parameter are:
// - bName
// - blaname=bName
// - WebSphere:blaname=bName
String blaID = "MyBLA"; // Replace the MyBLA value with your
                        // blaID value.
cmd.setParameter("blaID", blaID);

System.out.println("\nSet blaID parameter to "
    + cmd.getParameter("blaID"));

// Optionally set the cuID parameter.
String cuID = "myCU.zip"; // Replace the myCU.zip value with your
                          // cuID value.
cmd.setParameter("cuID", cuID);

System.out.println("\nSet cuID parameter to "
    + cmd.getParameter("cuID"));

// Optionally set the targetID parameter.
// The format of the targetID parameter for a cluster
// is WebSphere:cluster=cluster1
String targetID = "WebSphere:node=node1,server=server1"; // Replace
                                                         // this with your targetID value.
cmd.setParameter("targetID", targetID);

System.out.println("\nSet targetID parameter to "
    + cmd.getParameter("targetID"));

// Call the asynchronous client helper to process parameters.
try {
    asyncCmdClientHelper.processCommandParameters(cmd);
    System.out.println("\nCompleted processCommandParameters");
} catch (Throwable th) {
    System.out.println("Throwing an exception from " +
        "asyncCmdClientHelper.processCommandParameters(cmd).");
    th.printStackTrace();
    System.exit(-1);
}

// Run the command to check the status of the
// business-level application.
asyncCmdClientHelper.execute(cmd);
System.out.println("\nCompleted command execution");

CommandResult result = cmd.getCommandResult();
if (result != null) {
    if (result.isSuccessful()) {
        System.out.println("\nCommand executed successfully "
            + "with result\n" + result.getResult());
    }
    else {
        System.out.println("\nCommand executed with " +
            "Exception");
        result.getException().printStackTrace();
    }
}
} catch (Exception e) {
    e.printStackTrace();
}
}

package com.ibm.ws.management.application.task;

import com.ibm.websphere.management.cmdframework.provider.CommandNotification;
import com.ibm.websphere.management.async.client.AsyncCommandHandlerIF;

public class AsyncCmdTaskHandler implements AsyncCommandHandlerIF {

    public void handleNotification(CommandNotification notification) {

```



```
        // Add your own code here to handle the received notification.
        System.out.println("\nEXAMPLE: notification received: " +
            notification);
    }
}
```

What to do next

You can use the results of the status check to perform other tasks. For instance, if the results indicate that none of the composition units is started, you could start the business-level application.

Listing business-level applications using programming

You can list the business-level applications of a session so that you can complete further business-level application administration such as deleting a business-level application. A business-level application is an administrative model that captures the definition of an enterprise-level application so that you can perform specific business functions, such as accounting.

Before you begin

Before you can list business-level applications of a session, you must have created an empty business-level application.

About this task

You can list business-level applications of a session using programming, the administrative console, or the wsadmin tool. This topic describes how to list business-level applications using programming.

List all the business-level applications of a session unless you set the `blatID` parameter to specify the business-level application that you want to list. You can optionally list the business-level applications with a description for those that have a description if you set the `includeDescription` parameter to `true`. After you list the business-level applications, you can use the information to do further administration, such as starting or deleting business-level applications.

Perform the following tasks to list business-level applications of a session using programming.

Procedure

1. Connect to the application server.
The command framework allows the administrative command to be created and run with or without being connected to the application server. This step is optional if the application server is not running.
2. Create the command manager.
The command manager provides the functionality to create a new administrative command or query existing administrative commands.
3. Optionally create the asynchronous command handler for listening to command notifications.
Business-level application commands are implemented as asynchronous commands. To monitor the progress of the running command, you have to create an asynchronous command handler to receive notifications that the command generates.
4. Create the asynchronous command client.
An asynchronous command client provides a higher level interface to work with an asynchronous command. If you created an asynchronous command handler in the previous step, the handler is passed to the asynchronous command client. The asynchronous command client forwards the command notification to the handler and helps to control running of the command.
5. Use the command manager that you created in a previous step to create and set up the command that lists business-level applications of a session.

The command name is listBLAs. You can optionally set the blaID parameter to query for business-level applications that match the ID. You can optionally set the includeDescription parameter to display the business-level application descriptions.

6. Call the processCommandParameters method in the asynchronous command client to process the command parameters.

The command framework asynchronous command model requires this call.

7. Call the asynchronous command client to list the business-level applications of a session.

You might have created an asynchronous command handler to implement the AsyncCommandHandlerIF interface class in a previous step. If you did, the asynchronous command client listens to command notifications and forwards the notifications to the handler. The handler performs any necessary actions while waiting for the command to complete.

8. Check the command result when the command completes.

When the command finishes running, control is returned to the caller. You can then check the result by calling the command.getCommandResult method.

Results

After you successfully run the code, a list of business-level applications for a session is displayed.

Example

The following example shows how to list the business-level applications of a session based on the previous steps. Some statements are split on multiple lines for printing purposes.

```
package com.ibm.ws.management.application.task;

import java.util.Properties;

import com.ibm.websphere.management.AdminClient;
import com.ibm.websphere.management.AdminClientFactory;
import com.ibm.websphere.management.Session;
import com.ibm.websphere.management.cmdframework.AdminCommand;
import com.ibm.websphere.management.cmdframework.CommandMgr;
import com.ibm.websphere.management.cmdframework.CommandResult;
import com.ibm.websphere.management.async.client.AsyncCommandClient;

public class ListBLAs {

    public static void main(String[] args) {

        try {

            // Connect to the application server.
            // This step is optional if you use the local
            // command manager. Comment out the lines to and including
            // CommandMgr cmdMgr = CommandMgr.getClientCommandMgr(
            // soapClient);
            // to get the soapClient soap client if you use the local
            // command manager.

            String host = "localhost";
            String port = "8880"; // Change to your port number if
                                // it is not 8880.

            Properties config = new Properties();
            config.put(AdminClient.CONNECTOR_HOST, host);
            config.put(AdminClient.CONNECTOR_PORT, port);
            config.put(AdminClient.CONNECTOR_TYPE,
                    AdminClient.CONNECTOR_TYPE_SOAP);
            System.out.println("Config: " + config);
            AdminClient soapClient =
```

```

        AdminClientFactory.createAdminClient(config);

// Create the command manager.
CommandMgr cmdMgr = CommandMgr.getClientCommandMgr(soapClient);

// Comment out the previous lines to create a client command
// manager if you are using a local command manager.
// Uncomment the following line to create a local command
// manager:
//
// CommandMgr cmdMgr = CommandMgr.getCommandMgr();
System.out.println("\nCreated command manager");

// Optionally create an asynchronous command handler.
// Comment out the following line if no further handling
// of command notification is required:
AsyncCmdTaskHandler listener = new AsyncCmdTaskHandler();

// Create an asynchronous command client.

// Set up the session.
String id = Long.toHexString(System.currentTimeMillis());
String user = "content" + id;
Session session = new Session(user, true);

// If no command handler is used, replace listener with
// null for the AsyncCommandClient object.
AsyncCommandClient asyncCmdClientHelper = new
AsyncCommandClient(session, listener);
System.out.println("\nCreated async command client");

// Create the command that lists the business-level applications.
String cmdName = "listBLAs";
AdminCommand cmd = cmdMgr.createCommand(cmdName);
cmd.setConfigSession(session); // list all the business-level applications
// using the session created.

System.out.println("\nCreated " + cmdName);

// Optionally set the blaID parameter.
// Uncomment the following code to set the blaID parameter to
// only list the business-level applications with the ID specified. Otherwise all
// business-level applications are listed. Change the blaID parameter according
// to your scenario.
// Examples of valid formats for the blaID parameter are:
// - bName
// - blaname=bName
// - WebSphere:blaname=bName
// All business-level applications that match the ID specification
// are listed. The ID must include at least the business-level
// application name.
// String blaID = "bla1";
// cmd.setParameter("blaID", blaID);

//System.out.println("\nSet blaID parameter to "
// + cmd.getParameter("blaID"));

// Optionally include a description by setting
// the includeDescription parameter to true instead of false.
String includeDescription = "true";
cmd.setParameter("includeDescription", includeDescription);

System.out.println("\nSet includeDescription parameter to "
+ cmd.getParameter("includeDescription"));

// Call the asynchronous client helper to process parameters.

```

```

try {
    asyncCmdClientHelper.processCommandParameters(cmd);
    System.out.println("\nCompleted process command " +
        "parameters");
} catch (Throwable th) {
    System.out.println("Failed from " +
        "asyncCmdClientHelper.processCommandParameters(cmd).");
    th.printStackTrace();
    System.exit(-1);
}

// Run the command to list business-level applications.
asyncCmdClientHelper.execute(cmd);
System.out.println("\nCompleted running of the command");

// Check the command result.
CommandResult result = cmd.getCommandResult();
if (result != null) {
    if (result.isSuccessful()) {
        System.out.println("\nCommand ran successfully "
            + "with result\n" + result.getResult());
    }
    else {
        System.out.println("\nCommand ran with " +
            "Exception");
        result.getException().printStackTrace();
    }
}
} catch (Exception e) {
    e.printStackTrace();
}
}
}

```

```

package com.ibm.ws.management.application.task;

import com.ibm.websphere.management.cmdframework.provider.CommandNotification;
import com.ibm.websphere.management.async.client.AsyncCommandHandlerIF;

public class AsyncCmdTaskHandler implements AsyncCommandHandlerIF {

    public void handleNotification(CommandNotification notification) {
        // Add your own code here to handle the received notification
        System.out.println("\nEXAMPLE: notification received: " +
            notification);
    }
}

```

What to do next

You can complete other tasks associated with business-level applications, such as deleting, starting, or stopping business-level applications.

Listing composition units using programming

You can list the composition units for a specific business-level application so that you can complete further composition unit administration, such as deleting or adding composition units. A composition unit is typically created from a business-level application or an asset and contains configuration information that makes the asset runnable.

Before you begin

This task assumes a basic familiarity with command framework programming. Read about command framework programming in the application programming interfaces documentation.

Before you can list composition units, you must have imported an asset, created an empty business-level application, and added a composition unit to the business-level application.

About this task

You can list composition units using programming, the administrative console, or the wsadmin tool. This topic describes how to list composition units using programming.

You must provide the blaID parameter to specify the business-level application to list the composition unit. When you list composition units for a business-level application, you can optionally list the type for each composition unit and the description for each composition unit that has a description. You can use the list to complete further administration, such as deleting or exporting composition units.

Perform the following tasks to list composition units for a business-level application using programming.

Procedure

1. Connect to the application server.
The command framework allows the administrative command to be created and run with or without being connected to the application server. This step is optional if the application server is not running.
2. Create the command manager.
The command manager provides the functionality to create a new administrative command or query existing administrative commands.
3. Optionally create the asynchronous command handler for listening to command notifications.
Business-level application commands are implemented as asynchronous commands. To monitor the progress of the running command, you have to create an asynchronous command handler to receive notifications that the command generates.
4. Create the asynchronous command client.
An asynchronous command client provides a higher level interface to work with an asynchronous command. If you created an asynchronous command handler in the previous step, the handler is passed to the asynchronous command client. The asynchronous command client forwards the command notification to the handler and helps to control running of the command.
5. Use the command manager that you created in a previous step to create and set up the command that lists composition units.
The command name is listCompUnits. The blaID parameter is a required parameter that you use to specify the business-level application to list the composition units. You can optionally set the includeDescription parameter to display the composition unit descriptions. You can also optionally set the includeType parameter to display the composition unit types.
6. Call the processCommandParameters method in the asynchronous command client to process the command parameters.
The command framework asynchronous command model requires this call.
7. Call the asynchronous command client to list the composition units.
You might have created an asynchronous command handler to implement the AsyncCommandHandlerIF interface class in a previous step. If you did, the asynchronous command client listens to command notifications and forwards the notifications to the handler. The handler performs any necessary actions while waiting for the command to complete.
8. Check the command result when the command completes.

When the command finishes running, control is returned to the caller. You can then check the result by calling the `command.getCommandResult` method.

Results

After you successfully run the code, a list of composition units for a business-level application is displayed.

Example

The following example shows how to list the composition units of a specific business-level application based on the previous steps. Some statements are split on multiple lines for printing purposes.

```
package com.ibm.ws.management.application.task;

import java.util.Properties;

import com.ibm.websphere.management.AdminClient;
import com.ibm.websphere.management.AdminClientFactory;
import com.ibm.websphere.management.Session;
import com.ibm.websphere.management.cmdframework.AdminCommand;
import com.ibm.websphere.management.cmdframework.CommandMgr;
import com.ibm.websphere.management.cmdframework.CommandResult;
import com.ibm.websphere.management.async.client.AsyncCommandClient;

public class ListCompUnits {

    public static void main(String[] args) {

        try {

            // Connect to the application server.
            // This step is optional if you use the local
            // command manager. comment out the lines to and including
            // CommandMgr cmdMgr = CommandMgr.getClientCommandMgr(
            // soapClient);
            // to get the soapClient soap client if you use the local
            // command manager.

            String host = "localhost";
            String port = "8880"; // Change to your port number if
                // it is not 8880.

            Properties config = new Properties();
            config.put(AdminClient.CONNECTOR_HOST, host);
            config.put(AdminClient.CONNECTOR_PORT, port);
            config.put(AdminClient.CONNECTOR_TYPE,
                AdminClient.CONNECTOR_TYPE_SOAP);
            System.out.println("Config: " + config);
            AdminClient soapClient =
                AdminClientFactory.createAdminClient(config);

            // Create the command manager.
            CommandMgr cmdMgr = CommandMgr.getClientCommandMgr(soapClient);

            // Comment out the previous lines to create a client command
            // manager if you are using a local command manager.
            // Uncomment the following line to create a local command
            // manager:
            //
            // CommandMgr cmdMgr = CommandMgr.getCommandMgr();
            System.out.println("\nCreated command manager");

            // Optionally create an asynchronous command handler.
            // Comment out the following line if no further handling
            // of command notification is required:
            AsyncCmdTaskHandler listener = new AsyncCmdTaskHandler();
```

```

// Create an asynchronous command client.

// Set up the session.
String id = Long.toHexString(System.currentTimeMillis());
String user = "content" + id;
Session session = new Session(user, true);

// If no command handler is used, replace the listener with
// null for the AsyncCommandClient object.
AsyncCommandClient asyncCmdClientHelper = new
    AsyncCommandClient(session, listener);
System.out.println("\nCreated async command client");

// Create the command that lists the composition units.
String cmdName = "listCompUnits";
AdminCommand cmd = cmdMgr.createCommand(cmdName);
cmd.setConfigSession(session); // List all the composition units
                                // for the business-level application
                                // with this session ID.
System.out.println("\nCreated " + cmdName);

// Set the blaID parameter to the business-level application
// whose composition units are listing.
// Examples of valid formats for the blaID parameter are:
// - bName
// - blaname=bName
// - WebSphere:blaname=bName
// This parameter accepts an incomplete ID as long as the incomplete
// ID can resolve to a unique business-level application.
String blaID = "bla1";
cmd.setParameter("blaID", blaID);

System.out.println("\nSet blaID parameter to "
    + cmd.getParameter("blaID"));

// Optionally include descriptions for each composition unit
// that has a description by setting
// the includeDescription parameter to true or false.
String includeDescription = "true";
cmd.setParameter("includeDescription", includeDescription);

System.out.println("\nSet includeDescription parameter to "
    + cmd.getParameter("includeDescription"));
// Optionally include types for each composition unit
// by setting the includeType parameter to true or false.
String includeType = "true";
cmd.setParameter("includeType", includeType);

System.out.println("\nSet includeType parameter to "
    + cmd.getParameter("includeType"));

// Call the asynchronous client helper to process parameters.
try {
    asyncCmdClientHelper.processCommandParameters(cmd);
    System.out.println("\nCompleted process command " +
        "parameters");
} catch (Throwable th) {
    System.out.println("Failed from " +
        "asyncCmdClientHelper.processCommandParameters(cmd).");
    th.printStackTrace();
    System.exit(-1);
}

// Run the command to list the composition units.
asyncCmdClientHelper.execute(cmd);
System.out.println("\nCompleted running of command");

```

```

// Check the command result.
CommandResult result = cmd.getCommandResult();
if (result != null) {
    if (result.isSuccessful()) {
        System.out.println("\nCommand ran successfully "
            + "with result\n" + result.getResult());
    }
    else {
        System.out.println("\nCommand ran with " +
            "Exception");
        result.getException().printStackTrace();
    }
}
} catch (Exception e) {
    e.printStackTrace();
}
}
}

```

```
package com.ibm.ws.management.application.task;
```

```
import com.ibm.websphere.management.cmdframework.provider.CommandNotification;
import com.ibm.websphere.management.async.client.AsyncCommandHandlerIF;
```

```
public class AsyncCmdTaskHandler implements AsyncCommandHandlerIF {

    public void handleNotification(CommandNotification notification) {
        // Add your own code here to handle the received notification
        System.out.println("\nEXAMPLE: notification received: " +
            notification);
    }
}

```

What to do next

Now that you have listed the composition units for a business-level application, you can complete other tasks associated with composition units, such as adding or deleting composition units.

Listing control operations using programming

You can list the control operations of a business-level application or a composition unit for a session. You can use control operations, such as start or stop, to change or query the runtime environment of a business-level application or a composition unit.

Before you begin

Before you can list control operations of a business-level application or a composition unit for a session, you must have created an empty business-level application, imported an asset, and added a composition unit.

About this task

You can list control operations of a business-level application or a composition unit using programming, the administrative console, or the wsadmin tool. This topic describes how to list control operations using programming.

To list control operations for a business-level application of a session, provide a blaID parameter value, but no culD parameter value. To list control operations for a composition unit, specify both a blaID parameter value and a culD parameter value. To list all control operations for the specified business-level

application or the specified composition unit, do not specify an `opName` parameter value. To list the details for a specific control operation, set the `opName` parameter value to the name of the operation to list. To list details of the control operation definition, set the `long` parameter to `true`.

Perform the following tasks to list control operations for a business-level application or a composition unit of a session using programming.

Procedure

1. Connect to the application server.
The command framework allows the administrative command to be created and run with or without being connected to the application server. This step is optional if the application server is not running.
2. Create the command manager.
The command manager provides the functionality to create a new administrative command or query existing administrative commands.
3. Optionally create the asynchronous command handler for listening to command notifications.
Business-level application commands are implemented as asynchronous commands. To monitor the progress of the running command, you have to create an asynchronous command handler to receive notifications that the command generates.
4. Create the asynchronous command client.
An asynchronous command client provides a higher level interface to work with an asynchronous command. If you created an asynchronous command handler in the previous step, the handler is passed to the asynchronous command client. The asynchronous command client forwards the command notification to the handler and helps to control running of the command.
5. Create and set up the command that lists control operations of a business-level application or a composition unit of a session.
6. Call the `processCommandParameters` method in the asynchronous command client to process the command parameters.
The command framework asynchronous command model requires this call.
7. Call the asynchronous command client to list the control operations of a business-level application or a composition unit of a session.
You might have created an asynchronous command handler to implement the `AsyncCommandHandlerIF` interface class in a previous step. If you did, the asynchronous command client listens to command notifications and forwards the notifications to the handler. The handler performs any necessary actions while waiting for the command to complete.
8. Check the command result when the command completes.
When the command finishes running, control is returned to the caller. You can then check the result by calling the `command.getCommandResult` method.

Results

After you successfully run the code, a control operations of a business-level application or a composition unit for a session is displayed.

Example

The following example shows how to list the control operation of a business-level application or a composition unit of a session based on the previous steps. Some statements are split on multiple lines for printing purposes.

```
package com.ibm.ws.management.application.task;  
  
import java.util.Properties;  
  
import com.ibm.websphere.management.AdminClient;
```

```

import com.ibm.websphere.management.AdminClientFactory;
import com.ibm.websphere.management.Session;
import com.ibm.websphere.management.cmdframework.AdminCommand;
import com.ibm.websphere.management.cmdframework.CommandMgr;
import com.ibm.websphere.management.cmdframework.CommandResult;
import com.ibm.websphere.management.async.client.AsyncCommandClient;

public class listControlOps {

    public static void main(String[] args) {

        try {

            // Connect to the application server.
            // This step is optional if you use the local
            // command manager. Comment out the lines to and including
            // CommandMgr cmdMgr = CommandMgr.getClientCommandMgr(
            // soapClient);
            // to get the soapClient soap client if you use the local
            // command manager.
            String host = "localhost";
            String port = "8880"; // Change to your port number if
            // it is not 8880.

            Properties config = new Properties();
            config.put(AdminClient.CONNECTOR_HOST, host);
            config.put(AdminClient.CONNECTOR_PORT, port);
            config.put(AdminClient.CONNECTOR_TYPE,
                AdminClient.CONNECTOR_TYPE_SOAP);
            System.out.println("Config: " + config);
            AdminClient soapClient =
                AdminClientFactory.createAdminClient(config);

            // Create the command manager.
            CommandMgr cmdMgr = CommandMgr.getClientCommandMgr(soapClient);

            // Comment out the previous lines to create a client command
            // manager if you are using a local command manager.
            // Uncomment the following line to create a local command
            // manager:
            //
            // CommandMgr cmdMgr = CommandMgr.getCommandMgr();
            System.out.println("\nCreated command manager");

            // Optionally create an asynchronous command handler.
            // Comment out the following line if no further handling
            // of command notification is required:
            AsyncCmdTaskHandler listener = new AsyncCmdTaskHandler();

            // Create an asynchronous command client.

            // Set up the session.
            String id = Long.toHexString(System.currentTimeMillis());
            String user = "content" + id;
            Session session = new Session(user, true);

            // If no command handler is used, replace listener with
            // null for the AsyncCommandClient object.
            AsyncCommandClient asyncCmdClientHelper = new
                AsyncCommandClient(session, listener);
            System.out.println("\nCreated async command client");

            // Create the command that lists the control operations.
            String cmdName = "listControlOps";
            AdminCommand cmd = cmdMgr.createCommand(cmdName);
            cmd.setConfigSession(session); // List all the control operations
            // using the session created.

```

```

System.out.println("\nCreated " + cmdName);

// Set the blaID parameter, which is required.
// The blaID is for either the business-level application whose control
// units you are listing or for the business-level application whose
// composition unit control operations you are listing.
// Change the blaID parameter according to your
// scenario.
// Examples of valid formats for the blaID parameter are:
// - bName
// - blaname=bName
// - WebSphere:blaname=bName
// This parameter accepts an incomplete ID as long as the incomplete
// ID can resolve to a unique business-level application.
// String blaID = "bla1";
// cmd.setParameter("blaID", blaID);

// System.out.println("\nSet blaID parameter to "
//                     + cmd.getParameter("blaID"));

// Optionally set the cuID parameter to the composition
// unit whose control operations you are listing.
// Examples of valid formats for the cuID parameter are:
// - name
// - cuname=name
// - WebSphere:cuname=name
// This parameter accepts an incomplete ID as long as the
// incomplete ID can resolve to a unique composition unit
// within the business-level application.
//
// String cuID = "test5.zip";
// cmd.setParameter("cuID", cuID);

// System.out.println("\nSet cuID parameter to "
//                     + cmd.getParameter("cuID"));

// Optionally set the opName parameter of the operation to list.
// String opName = "opName1";
// cmd.setParameter("opName", opName);

// System.out.println("\nSet opnameID parameter to "
//                     + cmd.getParameter("opName"));

// Optionally include details of the control operation definition
// by setting the long parameter to true.
// String long = "true";
// cmd.setParameter("long", long);

// System.out.println("\nSet long parameter to "
//                     + cmd.getParameter("long"));

// Call the asynchronous client helper to process parameters.
try {
    asyncCmdClientHelper.processCommandParameters(cmd);
    System.out.println("\nCompleted process command " +
                      "parameters");
} catch (Throwable th) {
    System.out.println("Failed from " +
                      "asyncCmdClientHelper.processCommandParameters(cmd).");
    th.printStackTrace();
    System.exit(-1);
}

// Run the command to list control operations.
asyncCmdClientHelper.execute(cmd);

```

```

        System.out.println("\nCompleted running of command");

        // Check the command result.
        CommandResult result = cmd.getCommandResult();
        if (result != null) {
            if (result.isSuccessful()) {
                System.out.println("\nCommand ran successfully "
                    + "with result\n" + result.getResult());
            }
            else {
                System.out.println("\nCommand ran with " +
                    "Exception");
                result.getException().printStackTrace();
            }
        }
    } catch (Exception e) {
        e.printStackTrace();
    }
}
}
}

```

```

package com.ibm.ws.management.application.task;

import com.ibm.websphere.management.cmdframework.provider.CommandNotification;
import com.ibm.websphere.management.async.client.AsyncCommandHandlerIF;

public class AsyncCmdTaskHandler implements AsyncCommandHandlerIF {

    public void handleNotification(CommandNotification notification) {
        // Add your own code here to handle the received notification
        System.out.println("\nEXAMPLE: notification received: " +
            notification);
    }
}

```

What to do next

You can complete other tasks associated with business-level applications and composition units, such as deleting, starting, or stopping business-level applications or adding or exporting a composition unit.

Viewing a business-level application using programming

You can view business-level application information such as the description so that you can do other tasks associated with the business-level application, such as editing the business-level application. A business-level application is an administrative model that captures the entire definition of an enterprise-level application.

Before you begin

This task assumes a basic familiarity with command framework programming. Read about command framework programming in the application programming interface documentation.

Before you can view a business-level application, you must have created the business-level application.

You can view a business-level application using programming, the administrative console, or the wsadmin tool.

About this task

You must provide the `blaiD` parameter to specify the business-level application that you are viewing.

Perform the following tasks to view a business-level application using programming.

Procedure

1. Connect to the application server.
The command framework allows the administrative command to be created and run with or without being connected to the application server. This step is optional if the application server is not running.
2. Create the command manager.
The command manager provides the functionality to create a new administrative command or query existing administrative commands.
3. Optionally create the asynchronous command handler for listening to command notifications.
Business-level application commands are implemented as asynchronous commands. To monitor the progress of the running command, you have to create an asynchronous command handler to receive notifications that the command generates.
4. Create the asynchronous command client.
An asynchronous command client provides a higher level interface to work with an asynchronous command. If you created an asynchronous command handler in the previous step, the handler is passed to the asynchronous command client. The asynchronous command client forwards the command notification to the handler and helps to control running of the command.
5. Use the command manager that you created in a previous step to create and set up the command to view a business-level application.
The command name is `viewBLA`. Use the required `blaiD` parameter to specify the business-level application that you are viewing.
6. Call the `processCommandParameters` method in the asynchronous command client to process the command parameters.
The command framework asynchronous command model requires this call.
7. Display the command step.
8. Call the asynchronous command client to run the command to view a business-level application.
You could have created an asynchronous command handler to implement the `AsyncCommandHandlerIF` interface class in a previous step. If you did, the asynchronous command client listens to command notifications and forwards the notifications to the handler. The handler performs any necessary actions while waiting for the command to complete.
9. Check the command result when the command completes.
When the command finishes running, control is returned to the caller. You can then check the result by calling the `command.getCommandResult` method.

Results

After you successfully run the code, you can view a business-level application.

Example

The following example shows how to view a business-level application based on the previous steps. Some statements are split on multiple lines for printing purposes.

```
package com.ibm.ws.management.application.task;

import java.util.Properties;

import com.ibm.websphere.management.AdminClient;
```

```

import com.ibm.websphere.management.AdminClientFactory;
import com.ibm.websphere.management.Session;
import com.ibm.websphere.management.cmdframework.AdminCommand;
import com.ibm.websphere.management.cmdframework.CommandMgr;
import com.ibm.websphere.management.cmdframework.CommandResult;
import com.ibm.websphere.management.cmdframework.CommandStep;
import com.ibm.websphere.management.cmdframework.TaskCommand;
import com.ibm.websphere.management.async.client.AsyncCommandClient;

public class EditBLA {

    public static void main(String[] args) {

        try {

            // Connect to the application server.
            // This step is optional if you use the local command manager.
            // Comment out the following lines to get the soapClient soap client if
            // you are going to use the local command manager. You would
            // comment out the lines to and including
            // CommandMgr cmdMgr =
            // CommandMgr.getClientCommandMgr(soapClient);

            String host = "localhost"; // Change to your host if it is not localhost.
            String port = "8880"; // Change to your port number if it is not 8880.

            Properties config = new Properties();
            config.put(AdminClient.CONNECTOR_HOST, host);
            config.put(AdminClient.CONNECTOR_PORT, port);
            config.put(AdminClient.CONNECTOR_TYPE,
                AdminClient.CONNECTOR_TYPE_SOAP);
            System.out.println("Config: " + config);
            AdminClient soapClient =
                AdminClientFactory.createAdminClient(config);

            // Create the command manager.
            CommandMgr cmdMgr = CommandMgr.getClientCommandMgr(soapClient);

            // Comment out the previous lines to create a client command
            // manager if you are using a local command manager.
            // Uncomment the following line to create a local command
            // manager.
            //
            // CommandMgr cmdMgr = CommandMgr.getCommandMgr();

            System.out.println("\nCreated command manager");

            // Optionally create an asynchronous command handler.
            // Comment out the following line if no further handling
            // of command notification is required.
            AsyncCmdTaskHandler listener = new AsyncCmdTaskHandler();

            // Create an asynchronous command client.

            // Set up the session.
            // This example creates a new session. You can replace the
            // code below to use an existing session that has been
            // created.
            String id = Long.toHexString(System.currentTimeMillis());
            String user = "content" + id;
            Session session = new Session(user, true);

            // If no command handler is used, replace the listener with
            // null for the following AsyncCommandClient object.
            AsyncCommandClient asyncCmdClientHelper = new
                AsyncCommandClient(session, listener);
            System.out.println("\nCreated async command client");

```

```

// Create the command.
String cmdName = "viewBLA";
AdminCommand cmd = cmdMgr.createCommand(cmdName);
cmd.setConfigSession(session); // View an existing
                                // business-level application
                                // using the session created.
System.out.println("\nCreated " + cmdName);

// Set the required blaID parameter.
// Examples of valid formats for the blaID parameter are:
// - bName
// - blaname=bName
// - WebSphere:blaname=bName
// This parameter accepts an incomplete ID as long as the incomplete
// ID can resolve to a unique business-level application.
String blaID = "bla1"; // Replace the bla1 value with your value.
cmd.setParameter("blaID", blaID);

System.out.println("\nSet blaID parameter to "
    + cmd.getParameter("blaID"));

// Call the asynchronous client helper to process parameters.
try {
    asyncCmdClientHelper.processCommandParameters(cmd);
    System.out.println("\nCompleted process command " +
        "parameters");
} catch (Throwable th) {
    System.out.println("Throwing an exception from " +
        "asyncCmdClientHelper.processCommandParameters(cmd).");
    th.printStackTrace();
    System.exit(-1);
}

// Display step data.
String[] stepNames = ((TaskCommand) cmd).listCommandSteps();
for (int i = 0; i < stepNames.length; i++) {

    // Get the step.
    CommandStep step =
        ((TaskCommand)cmd).gotoStep(stepNames[i]);

    List paramNames = step.listParameterName();

    System.out.println("----- Step: " + step.getName() +
        " -----");
    // Get the parameter values for each row.
    for (int j = 0; j < step.getNumberOfRows(); j++) {
        System.out.println("  Row " + j);

        for (int k = 0; k < paramNames.size(); k++)
            System.out.println("      " + paramNames.get(k) +
                ": " + step.getParameter(
                    (String) paramNames.get(k), j));
    }
}

// Run the command to view the business-level application.
asyncCmdClientHelper.execute(cmd);
System.out.println("\nCompleted command execution");

CommandResult result = cmd.getCommandResult();
if (result != null) {
    if (result.isSuccessful()) {

```

```

        System.out.println("\nCommand executed successfully "
            + "with result\n" + result.getResult());
    }
    else {
        System.out.println("\nCommand executed with " +
            "Exception");
        result.getException().printStackTrace();
    }
}
} catch (Exception e) {
    e.printStackTrace();
}
}
}

package com.ibm.ws.management.application.task;

import com.ibm.websphere.management.cmdframework.provider.CommandNotification;
import com.ibm.websphere.management.async.client.AsyncCommandHandlerIF;

public class AsyncCmdTaskHandler implements AsyncCommandHandlerIF {

    public void handleNotification(CommandNotification notification) {
        // Add your own code here to handle the received notification.
        System.out.println("\nEXAMPLE: notification received: " +
            notification);
    }
}
}

```

What to do next

You can use the information that you viewed about the business-level application to perform other tasks. You might edit the business-level application to make improvements to it. You might start and stop a business-level application, delete a business-level application, add a composition unit to a business-level application, and so on.

Viewing a composition unit using programming

A composition unit is typically created from a business-level application or an asset and contains configuration information that makes the asset runnable. You can view the composition unit information so that you can complete other tasks associated with the composition unit such as editing an asset or delete a composition unit.

Before you begin

This task assumes a basic familiarity with command framework programming. Read about command framework programming in the application programming interfaces documentation.

Before you can view a composition unit of a business-level application, you must have created an empty business-level application, imported an asset into the business-level application, and added a composition unit to the business-level application.

About this task

You can view a composition unit using programming, the administrative console, or the wsadmin tool. This topic describes how to view a composition unit using programming.

You must provide the blaID and culID parameters to specify the composition unit of the business-level application that you are viewing. You can view configuration information of the composition unit of a

business-level application. The configuration information identifies the asset from which the composition unit is created if the composition unit contains an asset. You can also view runtime targets on which the deployable units of the composition unit are to run.

Perform the following tasks to view a composition unit of a business-level application using programming.

Procedure

1. Connect to the application server.
The command framework allows the administrative command to be created and run with or without being connected to the application server. This step is optional if the application server is not running.
2. Create the command manager.
The command manager provides the functionality to create a new administrative command or query existing administrative commands.
3. Optionally create the asynchronous command handler for listening to command notifications.
Business-level application commands are implemented as asynchronous commands. To monitor the progress of the running command, you have to create an asynchronous command handler to receive notifications that the command generates.
4. Create the asynchronous command client.
An asynchronous command client provides a higher level interface to work with an asynchronous command. If you created an asynchronous command handler in the previous step, the handler is passed to the asynchronous command client. The asynchronous command client forwards the command notification to the handler and helps to control running of the command.
5. Use the command manager that you created in a previous step to create and set up the command to view a composition unit.
The command name is `viewCompUnit`. Use the required `blalD` and `culD` parameters to specify the composition unit of the business-level application that you are viewing.
6. Call the `processCommandParameters` method in the asynchronous command client to process the command parameters.
The command framework asynchronous command model requires this call.
7. Call the asynchronous command client to run the command and view a composition unit.
You might have created an asynchronous command handler to implement the `AsyncCommandHandlerIF` interface class in a previous step. If you did, the asynchronous command client listens to command notifications and forwards the notifications to the handler. The handler performs any necessary actions while waiting for the command to complete.
8. Check the command result when the command completes.
When the command finishes running, control is returned to the caller. You can then check the result by calling the `command.getCommandResult` method.

Results

After you successfully run the code, you can view the configuration information of a composition unit for a business-level application.

Example

The following example shows how to view a composition unit of a business-level application based on the previous steps. Some statements are split on multiple lines for printing purposes.

```
package com.ibm.ws.management.application.task;

import java.util.Properties;

import com.ibm.websphere.management.AdminClient;
import com.ibm.websphere.management.AdminClientFactory;
```

```

import com.ibm.websphere.management.Session;
import com.ibm.websphere.management.cmdframework.AdminCommand;
import com.ibm.websphere.management.cmdframework.CommandMgr;
import com.ibm.websphere.management.cmdframework.CommandResult;
import com.ibm.websphere.management.cmdframework.CommandStep;
import com.ibm.websphere.management.cmdframework.TaskCommand;
import com.ibm.websphere.management.async.client.AsyncCommandClient;

public class ViewCompUnit {

    public static void main(String [] args) {

        try {
            // Connect to the application server.
            // This step is optional if you use the local
            // command manager. Comment out the lines to and including
            // CommandMgr cmdMgr = CommandMgr.getClientCommandMgr(
            // soapClient);
            // to get the soapClient soap client if you use the local
            // command manager.
            String host = "localhost";
            String port = "8880"; // Change to your port number if it is
            // not 8880.

            Properties config = new Properties();
            config.put(AdminClient.CONNECTOR_HOST, host);
            config.put(AdminClient.CONNECTOR_PORT, port);
            config.put(AdminClient.CONNECTOR_TYPE,
                AdminClient.CONNECTOR_TYPE_SOAP);
            System.out.println("Config: " + config);
            AdminClient soapClient =
                AdminClientFactory.createAdminClient(config);

            // Create the command manager.
            CommandMgr cmdMgr = CommandMgr.getClientCommandMgr(soapClient);

            // Comment out the previous lines to create a client command
            // manager if you are using a local command manager.
            // Uncomment the following line to create a local command
            // manager:
            //
            // CommandMgr cmdMgr = CommandMgr.getCommandMgr();
            System.out.println("\nCreated command manager");

            // Optionally create an asynchronous command handler.
            // Comment out the following line if no further handling
            // of command notification is required:
            AsyncCmdTaskHandler listener = new AsyncCmdTaskHandler();

            // Create an asynchronous command client.

            // Set up the session.
            String id = Long.toHexString(System.currentTimeMillis());
            String user = "content" + id;
            Session session = new Session(user, true);

            // If no command handler is used, replace the following listener with
            // null for the AsyncCommandClient object.
            AsyncCommandClient asyncCmdClientHelper = new
            AsyncCommandClient(session, listener);
            System.out.println("\nCreated async command client");

            // Create the command that views the composition unit.
            String cmdName = "viewCompUnit";
            AdminCommand cmd = cmdMgr.createCommand(cmdName);
            cmd.setConfigSession(session); // View a certain composition
            // unit of a business-level

```

```

        // application using the session created.
System.out.println("\nCreated " + cmdName);

// (required) Set the blaID parameter.
// Examples of valid formats for the blaID parameter are:
// - bName
// - blaname=bName
// - WebSphere:blaname=bName
// This parameter accepts an incomplete ID as long as the incomplete
// ID can resolve to a unique business-level application.
String blaID = "bla1";
cmd.setParameter("blaID", blaID);

System.out.println("\nSet blaID parameter to "
    + cmd.getParameter("blaID"));

// (required) Set the cuID parameter to the composition unit.
// The cuID parameter has the format of
// WebSphere:cuname=name. This parameter
// accepts an incomplete ID as long as the incomplete
// ID can resolve to a unique composition unit within the
// business-level application.
String cuID = "cu1";
cmd.setParameter("cuID", cuID);

System.out.println("\nSet cuID parameter to "
    + cmd.getParameter("cuID"));

// Call the asynchronous client helper to process parameters
try {
    asyncCmdClientHelper.processCommandParameters(cmd);
    System.out.println("\nCompleted process command " +
        "parameters");
} catch (Throwable th) {
    System.out.println("Failed from " +
        "asyncCmdClientHelper.processCommandParameters(cmd).");
    th.printStackTrace();
    System.exit(-1);
}

// Call the asynchronous command client to run the command.
asyncCmdClientHelper.execute(cmd);
System.out.println("\nCompleted running of the command");

// Check the command result.
CommandResult result = cmd.getCommandResult();
if (result != null) {
    if (result.isSuccessful()) {
        System.out.println("\nCommand ran successfully "
            + "with result\n" + result.getResult());
    }
    else {
        System.out.println("\nCommand ran with " +
            "Exception");
        result.getException().printStackTrace();
    }
}
} catch (Exception e) {
    e.printStackTrace();
}
}

package com.ibm.ws.management.application.task;

import com.ibm.websphere.management.cmdframework.provider.CommandNotification;
import com.ibm.websphere.management.async.client.AsyncCommandHandlerIF;

```

```
public class AsyncCmdTaskHandler implements AsyncCommandHandlerIF {
    public void handleNotification(CommandNotification notification) {
        // Add your own code here to handle the received notification
        System.out.println("\nEXAMPLE: notification received: " +
            notification);
    }
}
```

What to do next

You can use the information that you viewed about the composition unit to perform other tasks. For instance, you might edit the asset in the composition unit to make improvements to the asset. You might export the composition unit, and then import that composition unit into another business-level application.

Adding a composition unit using programming

You can add an asset to a business-level application by creating a composition unit for the asset. A composition unit is typically created from a business-level application or an asset and contains configuration information that makes the asset runnable.

Before you begin

Before you can add a composition unit to a business-level application, you must have created an empty business-level application and imported an asset.

You can add a composition unit to a business-level application using programming, the administrative console, or the wsadmin tool.

About this task

When you add a composition to a business-level application, the composition unit is configured for the specified business-level application. The composition unit cannot be shared with other business-level applications.

Perform the following steps to add a composition unit to a business-level application using programming.

Procedure

1. Connect to the application server.
The command framework allows the administrative command to be created and run with or without being connected to the application server. This step is optional if the application server is not running.
2. Create the command manager.
The command manager provides the functionality to create a new administrative command or query existing administrative commands.
3. Optionally create the asynchronous command handler for listening to command notifications.
Business-level application commands are implemented as asynchronous commands. To monitor the progress of the running command, you have to create an asynchronous command handler to receive notifications that the command generates.
4. Create the asynchronous command client.
An asynchronous command client provides a higher level interface to work with an asynchronous command. If you created an asynchronous command handler in the previous step, the handler is passed to the asynchronous command client. The asynchronous command client forwards the command notification to the handler and helps to control running of the command.
5. Use the command manager that you created in a previous step to create and set up the command that adds a composition unit.

The command name is `addCompUnit`. The `blaid` and `cuSourceID` parameters are required parameters that you use to specify composition unit source to be added to the business-level application. Examples of composition unit source are an asset or a business-level application. You can optionally provide deployable units for the composition unit through the `deplUnit` parameter. If the `cuSourceID` parameter is a Java Platform, Enterprise Edition (Java EE) asset, you can optionally use the `cuConfigStrategyFile` parameter or the `defaultBindingOptions` parameter to specify the default bindings. The `defaultBindingOptions` parameter must match the binding options available for this Java EE asset. To view a list of binding options available for this Java EE asset, look at the `AssetOptions` step in the `viewAsset` command. Specify each binding option in an `option_name=option_value` pair, with multiple pairs separated by a `#` character.

6. Call the `processCommandParameters` method in the asynchronous command client to process the command parameters.

The command framework asynchronous command model requires this call.

7. Set up the command step parameters.

You can set up composition unit information through various steps. The `CUOptions` step contains data about the composition unit such as its description, starting weight, and start and restart behavior. The `MapTargets` step contains target information about where the composition unit is to be deployed. The `RelationshipOptions` step contains shared library composition units on which this composition unit has dependencies. The `ActivationPlanOptions` step allows you to specify runtime components for each deployable unit. The `CreateAuxCUOptions` step contains assets on which this composition unit has dependencies. You can set up parameters in these steps.

8. Call the asynchronous command client to run the command that adds a composition unit to a business-level application.

You might have created an asynchronous command handler to implement the `AsyncCommandHandlerIF` interface class in a previous step. If you did, the asynchronous command client listens to command notifications and forwards the notifications to the handler. The handler performs any necessary actions while waiting for the command to complete.

9. Check the command result when the command completes.

When the command finishes running, control is returned to the caller. You can then check the result by calling the `command.getCommandResult` method.

Results

After you successfully run the code, the composition unit is added to the business-level application.

Example

The following example shows how to import an asset based on the previous steps. Some statements are split on multiple lines for printing purposes.

```
package com.ibm.ws.management.application.task;

import java.util.Properties;

import com.ibm.websphere.management.AdminClient;
import com.ibm.websphere.management.AdminClientFactory;
import com.ibm.websphere.management.Session;
import com.ibm.websphere.management.cmdframework.AdminCommand;
import com.ibm.websphere.management.cmdframework.CommandMgr;
import com.ibm.websphere.management.cmdframework.CommandResult;
import com.ibm.websphere.management.cmdframework.CommandStep;
import com.ibm.websphere.management.cmdframework.TaskCommand;
import com.ibm.websphere.management.async.client.AsyncCommandClient;

public class AddCompUnit {

    public static void main(String [] args) {
```

```

try {

    // Connect to the application server.
    // This step is optional if you use the local command
    // manager. Comment out the lines to and including
    // CommandMgr cmdMgr = CommandMgr.getClientCommandMgr(
    // soapClient);
    // to get the soapClient soap client if you use the
    // local command manager.

    String host = "localhost";
    String port = "8880"; // Change to your port number if it is
                        // not 8880.

    Properties config = new Properties();
    config.put(AdminClient.CONNECTOR_HOST, host);
    config.put(AdminClient.CONNECTOR_PORT, port);
    config.put(AdminClient.CONNECTOR_TYPE,
              AdminClient.CONNECTOR_TYPE_SOAP);
    System.out.println("Config: " + config);
    AdminClient soapClient =
        AdminClientFactory.createAdminClient(config);

    // Create the command manager
    CommandMgr cmdMgr = CommandMgr.getClientCommandMgr(soapClient);

    // Comment out the previous lines to create a client command
    // manager if you are using a local command manager.
    // Uncomment the following line to create a local command
    // manager:
    //
    // CommandMgr cmdMgr = CommandMgr.getCommandMgr();

    System.out.println("\nCreated command manager");

    // Optionally create the asynchronous command handler.
    // Comment out the following line if no further handling
    // of command notification is required:
    AsyncCmdTaskHandler listener = new AsyncCmdTaskHandler();

    // Create an asynchronous command client.

    // Set up the session.
    String id = Long.toHexString(System.currentTimeMillis());
    String user = "content" + id;
    Session session = new Session(user, true);

    // If no command handler is used, replace the following listener with
    // null for the AsyncCommandClient object.
    AsyncCommandClient asyncCmdClientHelper = new
    AsyncCommandClient(session, listener);
    System.out.println("\nCreated async command client");

    // Create the command to add a composition unit to a business-level application.
    String cmdName = "addCompUnit";
    AdminCommand cmd = cmdMgr.createCommand(cmdName);
    cmd.setConfigSession(session); // Add the composition unit using
                                // the session created.
    System.out.println("\nCreated " + cmdName);

    // Set the blaID command parameter.
    // Examples of valid formats for the blaID parameter are:
    // - bName
    // - blaname=bName
    // - WebSphere:blaname=bName
    // This parameter accepts an
    // incomplete ID as long as the incomplete

```

```

// ID can resolve to a unique business-level application.
String blaID = "bla1";
cmd.setParameter("blaID", blaID);

System.out.println("\nSet blaID parameter to "
    + cmd.getParameter("blaID"));

// Set the cuSourceID command parameter.
// Examples of valid formats for the cuSourceID parameter:
// If the source is an asset, examples are:
// - aName
// - assetname=aName
// - WebSphere:assetname=aName
// If the source is another business-level application,
// examples are:
// - bName
// - blaname=bName
// - WebSphere:blaname=bName
// The cuSourceID command parameter
// accepts an incomplete ID as long as the incomplete
// ID can resolve to a unique asset or business-level application.
String cuSourceID = "assetname=asset1.zip";
cmd.setParameter("cuSourceID", cuSourceID);

System.out.println("\nSet cuSourceID parameter to "
    + cmd.getParameter("cuSourceID"));

// Set the deplUnits command parameter.
// If the deployable units of an asset are, for example, a.jar and
// b.jar, then when you run the addCompUnit command you can
// specify deplUnits as a.jar+b.jar. You can specify the whole
// list, a subset of that list, or "default" to create this composition
// unit as a shared library. If the deplUnits parameter is not specified,
// the deployable units are set the same as that of their asset.
String deplUnits = "default";
cmd.setParameter("deplUnits", deplUnits);

System.out.println("\nSet deplUnits parameter to "
    + cmd.getParameter("deplUnits"));

// Call the asynchronous client helper to process parameters.
try {
    asyncCmdClientHelper.processCommandParameters(cmd);
    System.out.println("\nCompleted process command " +
        "parameters");
} catch (Throwable th) {
    System.out.println("Failed from " +
        "asyncCmdClientHelper.processCommandParameters(cmd).");
    th.printStackTrace();
    System.exit(-1);
}

// Set up the step parameters for the CUOptions step.
// The CUOptions step contains the following arguments:
// description - description for the composition unit
// startingWeight - starting weight for the composition
// unit within the business-level application. The default is 1.
// startedOnDistributed - to start composition unit upon distribution
// to target nodes. The default is false.
// restartBehaviorOnUpdate - restart behavior for a composition unit when
// updating the composition unit.
// The default is DEFAULT. Valid values are DEFAULT, ALL, and NONE.
String stepName = "CUOptions";
CommandStep step = ((TaskCommand) cmd).gotoStep(stepName);

// Composition unit name:
String name = "cu1";

```

```

// Composition unit description:
String description = "cul description";

for(int i = 0; i < step.getNumberOfRows(); i++) {
    // The following lines change the composition unit name and
    // description step parameters of the CUOptions step. Change
    // your set of step parameters as required for your
    // scenario.

    // Set the name.
    step.setParameter("name", name, i);
    System.out.println("\nSet name parameter to " +
        step.getParameter("name", i));

    // Set the description.
    step.setParameter("description", description, i);
    System.out.println("\nSet description parameter to " +
        step.getParameter("description", i));
}

// Set up the step parameters for the MapTargets step.
stepName = "MapTargets";
step = ((TaskCommand) cmd).gotoStep(stepName);

// Specify the targets to deploy the composition unit.
// The default is server1. Use the + character to
// specify multiple targets.
String server = "server1";

for(int i = 0; i < step.getNumberOfRows(); i++) {
    // The following lines change the composition unit and
    // server step parameters of the
    // MapTargets step. Change your set of step parameters
    // as required for your scenario.

    // Set the server.
    step.setParameter("server", server, i);
    System.out.println("\nSet server parameter to " +
        step.getParameter("server", i));
}

// The addCompUnit command might contain the
// CreateAuxCUOptions, RelationshipOptions and ActivationPlanOptions
// steps, depending on the asset content of the assets imported.
// The CreateAuxCUOptions step is available if the cuSourceID value
// is an asset. The asset includes an asset relationship to an
// asset that does not have a matching composition unit in the
// business-level application.
//
// If the CreateAuxCUOptions step is available, the selected
// deployable units of the source asset of the "primary" composition
// unit (that is, the composition unit being added) have dependencies
// on other assets for which there are no matching composition units
// in the business-level application. A "secondary" composition unit will be created for each
// of those asset dependencies.
//
// Each CreateAuxCUOptions row corresponds to one dependency
// relationship declaration. Each row consists of parameter values
// for the dependency relationship. Some parameters are read-only and
// some of them are editable. To edit parameter values, use the same
// approach as that used to edit parameter values in the CUOptions step.
//
// The parameters for this step include:
//
//   deplUnit - The name of the deployable unit which has the
//              dependency. (Read-only.)

```



```

// inputAsset - The asset ID for the source asset of the primary
// composition unit. (Read-only.)
// cuID - The name of the secondary composition unit to create.
// matchTarget - Specifies whether the server target for the secondary
// composition unit is to match the server target for
// the primary composition unit. The default value
// is "true". If the value is set to "false", the
// secondary composition unit will be created with no
// target. The target on the secondary composition unit
// can be set at a later time with the editCompUnit
// command.
//
// If the RelationshipOptions step is available, the selected
// deployable units of the source asset of the "primary" composition
// unit (that is, the composition unit being added) have dependencies
// on other assets for which there are matching "secondary" composition
// units in the business-level application. The RelationshipOptions step is much like
// CreateAuxCUOptions except that the required secondary composition
// units already exist. Also, each RelationshipOptions row maps one
// deployable unit to one or more secondary composition units, whereas,
// each CreateAuxCUOptions row maps one deployable unit to one
// asset dependency.
//
// Each RelationshipOptions row corresponds to one deployable unit
// with one or more dependency relationships and consists of
// parameter values for the dependency relationships. Some parameters
// are read-only and some of them are editable. To edit parameter
// values, use the same approach as that used to edit parameter values
// in the CUOptions step.
//
// The parameters for this step include:
//
// deplUnit - The name of the deployable unit which has the
// dependency. (Read-only.)
// relationship - Composition unit dependencies in the form of a
// list of composition unit IDs. Composition unit
// IDs are separated by a "plus" sign (+). Each ID
// can be fully or partially formed as shown with the
// following examples:
//     WebSphere:cuname=SharedLib1.jar
//     WebSphere:cuname=SharedLib.jar
//     SharedLib.jar
// matchTarget - Specifies whether the server target for the secondary
// composition units are to match the server target for
// the primary composition unit. The default value
// is "true". If the value is set to "false", the
// secondary composition unit will be created with no
// target. The target on the secondary composition unit
// can be set at a later time with the editCompUnit
// command.
//
// The addCompUnit command contains the ActivationPlanOptions step.
// The user can set the ActivationPlanOptions step parameters
// similar to the step parameters for the CUOptions step in
// the previous examples. The arguments for this step include:
// deplUnit - deployable unit URI (read only parameter)
// activationPlan - specifies a list of runtime components in the
// format of specname=xxxx
//
// Run the command to add the composition unit.
asyncCmdClientHelper.execute(cmd);
System.out.println("\nCompleted running of command");

// Check the command result.
CommandResult result = cmd.getCommandResult();
if (result != null) {
    if (result.isSuccessful()) {
        System.out.println("\nCommand ran successfully "

```

```

        + "with result\n" + result.getResult());
    }
    else {
        System.out.println("\nCommand ran with " +
            "Exception");
        result.getException().printStackTrace();
    }
}
} catch (Exception e) {
    e.printStackTrace();
}
}
}

package com.ibm.ws.management.application.task;

import com.ibm.websphere.management.cmdframework.provider.CommandNotification;
import com.ibm.websphere.management.async.client.AsyncCommandHandlerIF;

public class AsyncCmdTaskHandler implements AsyncCommandHandlerIF {

    public void handleNotification(CommandNotification notification) {
        // Add your own code here to handle the received notification
        System.out.println("\nEXAMPLE: notification received: " +
            notification);
    }
}

```

What to do next

Start the business-level application to which you added the composition unit. Complete administrative tasks such as viewing or deleting the composition unit.

Updating an asset using programming

You can update an asset by adding, deleting, or updating a single file or Java Platform, Enterprise Edition (Java EE) module, or by merging multiple files or Java EE modules into an asset. You can also update an asset by replacing the entire asset.

Before you begin

This task assumes a basic familiarity with command framework programming. Read about command framework programming in the application programming interface documentation.

Before you can update an asset, you must have imported the asset.

You can update an asset using programming, the administrative console, or the wsadmin tool.

About this task

You must specify the `assetID` parameter of the asset that you are updating. In addition, you must specify the operation parameter. Whether or not you must specify the `contents` and `contenturi` parameters depends on the operation that you specify.

You modify one or more files or module files of an asset with this task. You also update the asset binary file, but do not update the composition units that the system deploys with this asset as a backing object.

Perform the following tasks to update an asset using programming.

Procedure

1. Connect to the application server.
The command framework allows the administrative command to be created and run with or without being connected to the application server. This step is optional if the application server is not running.
2. Create the command manager.
The command manager provides the functionality to create a new administrative command or query existing administrative commands.
3. Optionally create the asynchronous command handler for listening to command notifications.
Business-level application commands are implemented as asynchronous commands. To monitor the progress of the running command, you have to create an asynchronous command handler to receive notifications that the command generates.
4. Create the asynchronous command client.
An asynchronous command client provides a higher level interface to work with an asynchronous command. If you created an asynchronous command handler in the previous step, the handler is passed to the asynchronous command client. The asynchronous command client forwards the command notification to the handler and helps to control running of the command.
5. Create and set up the command that updates an asset.
 - a. Set the parameter for the asset that you are updating.
 - b. Set the operation parameter.
 - c. Set the contents parameter unless the operation is set to delete.
 - d. Set the contenturi parameter if the operation is set to add, update, or addupdate.
6. Call the processCommandParameters method in the asynchronous command client to process the command parameters.
The command framework asynchronous command model requires this call.
7. Call the asynchronous command client to run the command to update an asset.
You could have created an asynchronous command handler to implement the AsyncCommandHandlerIF interface class in a previous step. If you did, the asynchronous command client listens to command notifications and forwards the notifications to the handler. The handler performs any necessary actions while waiting for the command to complete.
8. Check the command result when the command completes.
When the command finishes running, control is returned to the caller. You can then check the result by calling the command.getCommandResult method.

Results

After you successfully run the code, the asset is updated.

Example

The following example shows how to update an asset based on the previous steps. Some statements are split on multiple lines for printing purposes.

```
package com.ibm.ws.management.application.task;

import java.util.Properties;

import com.ibm.websphere.management.AdminClient;
import com.ibm.websphere.management.AdminClientFactory;
import com.ibm.websphere.management.Session;
import com.ibm.websphere.management.cmdframework.AdminCommand;
import com.ibm.websphere.management.cmdframework.CommandMgr;
import com.ibm.websphere.management.cmdframework.CommandResult;
import com.ibm.websphere.management.cmdframework.CommandStep;
import com.ibm.websphere.management.cmdframework.TaskCommand;
```

```

import com.ibm.websphere.management.async.client.AsyncCommandClient;

public class EditBLA {

    public static void main(String[] args) {

        try {

            // Connect to the application server.
            // This step is optional if you use the local command manager.
            // Comment out the following lines to get soapClient soap client if
            // you are going to use the local command manager.
            // Comment out the lines to and including
            // CommandMgr cmdMgr =
            // CommandMgr.getClientCommandMgr(soapClient);

            String host = "localhost"; // Change to your host if it is not localhost.
            String port = "8880"; // Change to your port number if it is not 8880.

            Properties config = new Properties();
            config.put(AdminClient.CONNECTOR_HOST, host);
            config.put(AdminClient.CONNECTOR_PORT, port);
            config.put(AdminClient.CONNECTOR_TYPE,
                AdminClient.CONNECTOR_TYPE_SOAP);
            System.out.println("Config: " + config);
            AdminClient soapClient =
                AdminClientFactory.createAdminClient(config);

            // Create the command manager.
            CommandMgr cmdMgr = CommandMgr.getClientCommandMgr(soapClient);

            // Comment out the previous lines to create a client command
            // manager if you are using a local command manager.
            // Uncomment the following line to create a local command
            // manager.
            //
            // CommandMgr cmdMgr = CommandMgr.getCommandMgr();

            System.out.println("\nCreated command manager");

            // Optionally create an asynchronous command handler.
            // Comment out the following line if no further handling
            // of command notification is required.
            AsyncCmdTaskHandler listener = new AsyncCmdTaskHandler();

            // Create an asynchronous command client.

            // Set up the session.
            // This example creates a new session. You can replace the
            // following code to use an existing session that has been
            // created.
            String id = Long.toHexString(System.currentTimeMillis());
            String user = "content" + id;
            Session session = new Session(user, true);

            // If you do not use the command handler, replace the listener with
            // null for the following AsyncCommandClient object.
            AsyncCommandClient asyncCmdClientHelper = new
                AsyncCommandClient(session, listener);
            System.out.println("\nCreated async command client");

            // Create the command that updates the asset.
            String cmdName = "updateAsset";
            AdminCommand cmd = cmdMgr.createCommand(cmdName);
            cmd.setConfigSession(session); // Update an asset
                                         // using the session
                                         // created.

```

```

System.out.println("\nCreated " + cmdName);

// Set the required assetID parameter.
// Examples of valid formats for the assetID parameter:
// - aName
// - assetname=aName
// - WebSphere:assetname=aName
// This parameter accepts an incomplete ID as long as the
// incomplete ID can resolve to a unique asset within the
// business-level application.
String assetID = "asset1.zip"; // Replace asset1.zip with your
                               // value of the assetID parameter.
cmd.setParameter("assetID", assetID);

System.out.println("\nSet assetID parameter to "
    + cmd.getParameter("assetID"));

// Set the required operation parameter.
// Possible operation values are add, addupdate, delete, merge,
// replace, and update.
// Use the add value to add a new file or Java EE module to the asset.
// Use the addupdate value to add a new file or Java EE module to the asset, or
// update an existing file or Java EE module.
// Use the delete value to delete an existing file or Java EE module in the asset.
// Use the merge value to provide a partial update with multiple
// additions, updates, or deletions.
// Use the replace value for a full update to replace all the contents.
// Use the update value to update an existing file or Java EE module in the asset.
String op = "add"; // Replace the add value with your operation value.
cmd.setParameter("operation", op);

System.out.println("\nSet operation parameter to "
    + cmd.getParameter("operation"));

// Set the contents parameter.
// This parameter is required unless the operation is set to
// delete.
String contents = "/assets/abc.txt"

cmd.setParameter("contents", contents);

System.out.println("\nSet contents parameter to "
    + cmd.getParameter("contents"));

// Set the contenturi parameter.
// This parameter is required for the
// add, addupdate, update, or delete operations.
String contenturi = "abc.txt"; // URI within the asset to
                               // place the new file. Replace
                               // with your value.
cmd.setParameter("contenturi", contenturi);

System.out.println("\nSet contenturi parameter to "
    + cmd.getParameter("contenturi"));

// Call the asynchronous client helper to process parameters.
try {
    asyncCmdClientHelper.processCommandParameters(cmd);
    System.out.println("\nCompleted process command " +
        "parameters");
} catch (Throwable th) {
    System.out.println("Throwing an exception from " +
        "asyncCmdClientHelper.processCommandParameters(cmd).");
    th.printStackTrace();
    System.exit(-1);
}

```

```

// Run the command.
asyncCmdClientHelper.execute(cmd);
System.out.println("\nCompleted command execution");

CommandResult result = cmd.getCommandResult();
if (result != null) {
    if (result.isSuccessful()) {
        System.out.println("\nCommand executed successfully "
            + "with result\n" + result.getResult());
    }
    else {
        System.out.println("\nCommand executed with " +
            "Exception");
        result.getException().printStackTrace();
    }
}
} catch (Exception e) {
    e.printStackTrace();
}
}
}

package com.ibm.ws.management.application.task;

import com.ibm.websphere.management.cmdframework.provider.CommandNotification;
import com.ibm.websphere.management.async.client.AsyncCommandHandlerIF;

public class AsyncCmdTaskHandler implements AsyncCommandHandlerIF {

    public void handleNotification(CommandNotification notification) {
        // Add your own code here to handle the received notification.
        System.out.println("\nEXAMPLE: notification received: " +
            notification);
    }
}
}

```

What to do next

You can do other tasks associated with assets in business-level applications, such as adding or deleting other assets, listing assets, exporting assets, and so on.

Editing a business-level application using programming

You can edit the information of a business-level application such as its description. A business-level application is an administrative model that captures the entire definition of an enterprise-level application.

Before you begin

This task assumes a basic familiarity with command framework programming. Read about command framework programming in the application programming interfaces documentation.

Before you can edit a business-level application, you must have created a business-level application.

You can edit a business-level application using programming, the administrative console, or the wsadmin tool.

About this task

You must provide the blaID parameter to specify the business-level application that you are editing.

Perform the following tasks to edit a business-level application using programming.

Procedure

1. Connect to the application server.
The command framework allows the administrative command to be created and run with or without being connected to the application server. This step is optional if the application server is not running.
2. Create the command manager.
The command manager provides the functionality to create a new administrative command or query existing administrative commands.
3. Optionally create the asynchronous command handler for listening to command notifications.
Business-level application commands are implemented as asynchronous commands. To monitor the progress of the running command, you have to create an asynchronous command handler to receive notifications that the command generates.
4. Create the asynchronous command client.
An asynchronous command client provides a higher level interface to work with an asynchronous command. If you created an asynchronous command handler in the previous step, the handler is passed to the asynchronous command client. The asynchronous command client forwards the command notification to the handler and helps to control running of the command.
5. Use the command manager that you created in a previous step to create the command that edits a business-level application.
The command name is editBLA. Use the required blaID parameter to specify the business-level application that you are editing.
6. Call the processCommandParameters method in the asynchronous command client to process the command parameters.
The command framework asynchronous command model requires this call.
7. Set up the command step parameter by setting the description parameter.
The BLAOptions step contains a description for the business-level application. You can edit the description parameter in the BLAOptions step.
8. Call the asynchronous command client to run the command to edit a business-level application.
You could have created an asynchronous command handler to implement the AsyncCommandHandlerIF interface class in a previous step. If you did, the asynchronous command client listens to command notifications and forwards the notifications to the handler. The handler performs any necessary actions while waiting for the command to complete.
9. Check the command result when the command completes.
When the command finishes running, control is returned to the caller. You can then check the result by calling the command.getCommandResult method.

Results

After you successfully run the code, the business-level application is edited.

Example

The following example shows how to edit a business-level application based on the previous steps. Some statements are split on multiple lines for printing purposes.

```
package com.ibm.ws.management.application.task;

import java.util.Properties;

import com.ibm.websphere.management.AdminClient;
import com.ibm.websphere.management.AdminClientFactory;
import com.ibm.websphere.management.Session;
import com.ibm.websphere.management.cmdframework.AdminCommand;
import com.ibm.websphere.management.cmdframework.CommandMgr;
import com.ibm.websphere.management.cmdframework.CommandResult;
```

```

import com.ibm.websphere.management.cmdframework.CommandStep;
import com.ibm.websphere.management.cmdframework.TaskCommand;
import com.ibm.websphere.management.async.client.AsyncCommandClient;

public class EditBLA {

    public static void main(String[] args) {
        try {

            // Connect to the application server.
            // This step is optional if you use the local command manager.
            // Comment out the following lines to get the soapClient SOAP client if
            // you are going to use the local command manager. You would
            // comment out the lines to and including
            // CommandMgr cmdMgr =
            // CommandMgr.getClientCommandMgr(soapClient);

            String host = "localhost"; // Change to your host if it is not localhost.
            String port = "8880"; // Change to your port number if it is not 8880.

            Properties config = new Properties();
            config.put(AdminClient.CONNECTOR_HOST, host);
            config.put(AdminClient.CONNECTOR_PORT, port);
            config.put(AdminClient.CONNECTOR_TYPE,
                AdminClient.CONNECTOR_TYPE_SOAP);
            System.out.println("Config: " + config);
            AdminClient soapClient =
                AdminClientFactory.createAdminClient(config);

            // Create command manager.
            CommandMgr cmdMgr = CommandMgr.getClientCommandMgr(soapClient);

            // Comment out the previous lines to create a client command
            // manager if you are using a local command manager.
            // Uncomment the following line to create a local command
            // manager.
            //
            // CommandMgr cmdMgr = CommandMgr.getCommandMgr();

            System.out.println("\nCreated command manager");

            // Optionally create an asynchronous command handler.
            // Comment out the following line if no further handling
            // of command notification is required.
            AsyncCmdTaskHandler listener = new AsyncCmdTaskHandler();

            // Create an asynchronous command client.

            // Set up the session.
            // This example creates a new session. You can replace the
            // code below to use an existing session that has been
            // created.
            String id = Long.toHexString(System.currentTimeMillis());
            String user = "content" + id;
            Session session = new Session(user, true);

            // If no command handler is used, replace the listener with
            // null for the following AsyncCommandClient object.
            AsyncCommandClient asyncCmdClientHelper = new
                AsyncCommandClient(session, listener);
            System.out.println("\nCreated async command client");

            // Create the command that edits the business-level application.
            String cmdName = "editBLA";
            AdminCommand cmd = cmdMgr.createCommand(cmdName);
            cmd.setConfigSession(session); // Edit an existing business-level
            // application using the session

```



```

        // created.
System.out.println("\nCreated " + cmdName);

// Set the blaID parameter (required).
// Examples of valid formats for the blaID parameter are:
// - bName
// - blaname=bName
// - WebSphere:blaname=bName
// This parameter accepts an incomplete ID as long as the incomplete
// ID can resolve to a unique business-level application.
String blaID = "bla1"; // Replace bla1 with your value of the blaID.
cmd.setParameter("blaID", blaID);

System.out.println("\nSet blaID parameter to "
    + cmd.getParameter("blaID"));

// Call the asynchronous client helper to process parameters.
try {
    asyncCmdClientHelper.processCommandParameters(cmd);
    System.out.println("\nCompleted process command " +
        "parameters");
} catch (Throwable th) {
    System.out.println("Throwing an exception from " +
        "asyncCmdClientHelper.processCommandParameters(cmd).");
    th.printStackTrace();
    System.exit(-1);
}

// Set up the step parameters for the BLAOptions step.
// The only step parameter you can edit is description.
String stepName = "BLAOptions";
CommandStep step = ((TaskCommand) cmd).gotoStep(stepName);

// Edit the business-level application description.
String description = "bla for testing"; // Replace with your value.

for (int i = 0; i < step.getNumberOfRows(); i++) {
    // The following lines set the description
    // step parameter.
    step.setParameter("description", description, i);
    System.out.println("\nSet description parameter to " +
        step.getParameter("description", i));
}

// Run the command to edit the business-level application.
asyncCmdClientHelper.execute(cmd);
System.out.println("\nCompleted command execution");

// Check the command result.
CommandResult result = cmd.getCommandResult();
if (result != null) {
    if (result.isSuccessful()) {
        System.out.println("\nCommand executed successfully "
            + "with result\n" + result.getResult());
    }
    else {
        System.out.println("\nCommand executed with " +
            "Exception");
        result.getException().printStackTrace();
    }
}
} catch (Exception e) {
    e.printStackTrace();
}
}
}

```

```

package com.ibm.ws.management.application.task;

import com.ibm.websphere.management.cmdframework.provider.CommandNotification;
import com.ibm.websphere.management.async.client.AsyncCommandHandlerIF;

public class AsyncCmdTaskHandler implements AsyncCommandHandlerIF {

    public void handleNotification(CommandNotification notification) {
        // Add your own code here to handle the received notification.
        System.out.println("\nEXAMPLE: notification received: " +
            notification);
    }
}

```

What to do next

After you edit the business-level application, you can continue administration of business-level applications. You can do such things as start and stop a business-level application, delete a business-level application, add a composition unit to a business-level application, and so on.

Editing a composition unit using programming

You can edit the configuration information in a composition unit of a business-level application if, for example, you want to change certain modules in the composition unit that are configured to run in specific targets. A composition unit is typically created from a business-level application or an asset and contains configuration information that makes the asset runnable.

Before you begin

This task assumes a basic familiarity with command framework programming. Read about command framework programming in the application programming interfaces documentation.

Before you can edit a composition unit of a business-level application, you must have created an empty business-level application, imported an asset, and added a composition unit to the business-level application.

About this task

You can edit a composition unit of a business-level application using programming, the administrative console, or the wsadmin tool. This topic describes how to edit a composition unit of a business-level application using programming.

You must provide the blaID and culD parameters to specify the composition unit of the business-level application that you are editing.

Perform the following tasks to edit a composition unit of a business-level application using programming.

Procedure

1. Connect to the application server.
The command framework allows the administrative command to be created and run with or without being connected to the application server. This step is optional if the application server is not running.
2. Create the command manager.
The command manager provides the functionality to create a new administrative command or query existing administrative commands.
3. Optionally create the asynchronous command handler for listening to command notifications.

Business-level application commands are implemented as asynchronous commands. To monitor the progress of the running command, you have to create an asynchronous command handler to receive notifications that the command generates.

4. Create the asynchronous command client.

An asynchronous command client provides a higher level interface to work with an asynchronous command. If you created an asynchronous command handler in the previous step, the handler is passed to the asynchronous command client. The asynchronous command client forwards the command notification to the handler and helps to control running of the command.

5. Use the command manager that you created in a previous step to create and set up the command that edits a composition unit of a business-level application.

The command name is `editCompUnit`. Use the required `blaiD` and `culD` parameters to specify the composition unit of the business-level application that you are editing.

6. Call the `processCommandParameters` method in the asynchronous command client to process the command parameters.

The command framework asynchronous command model requires this call.

7. Set up the command step parameters.

You can edit various composition unit information through steps. The `CUOptions` step contains data about the composition unit such as its description, starting weight, and start and restart behavior. The `MapTargets` step contains target information about where to deploy the composition unit. The `RelationshipOptions` step contains shared library composition units on which this composition unit has a dependency. The `ActivationPlanOptions` step allows you to change runtime components for each deployable unit. You can edit parameters in these steps.

8. Call the asynchronous command client to run the command that edits a composition unit of a business-level application.

You might have created an asynchronous command handler to implement the `AsyncCommandHandlerIF` interface class in a previous step. If you did, the asynchronous command client listens to command notifications and forwards the notifications to the handler. The handler performs any necessary actions while waiting for the command to complete.

9. Check the command result when the command completes.

When the command finishes running, control is returned to the caller. You can then check the result by calling the `command.getCommandResult` method.

Results

After you successfully run the code, the composition unit of a business-level application is edited.

Example

The following example shows how to edit a composition unit of a business-level application based on the previous steps. Some statements are split on multiple lines for printing purposes.

```
package com.ibm.ws.management.application.task;

import java.util.Properties;

import com.ibm.websphere.management.AdminClient;
import com.ibm.websphere.management.AdminClientFactory;
import com.ibm.websphere.management.Session;
import com.ibm.websphere.management.cmdframework.AdminCommand;
import com.ibm.websphere.management.cmdframework.CommandMgr;
import com.ibm.websphere.management.cmdframework.CommandResult;
import com.ibm.websphere.management.cmdframework.CommandStep;
import com.ibm.websphere.management.cmdframework.TaskCommand;
import com.ibm.websphere.management.async.client.AsyncCommandClient;

public class EditCompUnit {
```

```

public static void main(String [] args) {

    try {

        // Connect to the application server.
        // This step is optional if you use the local
        // command manager. Comment out the lines to and including
        // CommandMgr cmdMgr = CommandMgr.getClientCommandMgr(
        // soapClient);
        // to get the soapClient soap client if you use the local
        // command manager.

        String host = "localhost";
        String port = "8880"; // Change to your port number if it is
            // not 8880.

        Properties config = new Properties();
        config.put(AdminClient.CONNECTOR_HOST, host);
        config.put(AdminClient.CONNECTOR_PORT, port);
        config.put(AdminClient.CONNECTOR_TYPE,
            AdminClient.CONNECTOR_TYPE_SOAP);
        System.out.println("Config: " + config);
        AdminClient soapClient =
            AdminClientFactory.createAdminClient(config);

        // Create the command manager.
        CommandMgr cmdMgr = CommandMgr.getClientCommandMgr(soapClient);

        // Comment out the previous lines to create a client command
        // manager if you are using a local command manager.
        // Uncomment the following line to create a local command
        // manager:
        // CommandMgr cmdMgr = CommandMgr.getCommandMgr();
        System.out.println("\nCreated command manager");

        // Optionally create an asynchronous command handler.
        // Comment out the following line if no further handling
        // of command notification is required:
        AsyncCmdTaskHandler listener = new AsyncCmdTaskHandler();

        // Create an asynchronous command client.

        // Set up the session.
        String id = Long.toHexString(System.currentTimeMillis());
        String user = "content" + id;
        Session session = new Session(user, true);

        // If no command handler is used, replace the listener with
        // null for the following AsyncCommandClient object:
        AsyncCommandClient asyncCmdClientHelper = new
        AsyncCommandClient(session, listener);
        System.out.println("\nCreated async command client");

        // Create the command that edits the composition unit.
        String cmdName = "editCompUnit";
        AdminCommand cmd = cmdMgr.createCommand(cmdName);
        cmd.setConfigSession(session); // Edit a certain composition
            // unit of a business-level using the session created.
        System.out.println("\nCreated " + cmdName);

        // Set the blaID parameter.
        // Examples of valid formats for the blaID parameter are:
        // - bName
        // - blaname=bName
        // - WebSphere:blaname=bName
        // This parameter accepts an incomplete ID as long as the

```

```

// incomplete ID can resolve to a unique business-level application.
String blaID = "bla1";
cmd.setParameter("blaID", blaID);

System.out.println("\nSet blaID parameter to "
    + cmd.getParameter("blaID"));

// Set the cuID parameter.
// Examples of valid formats for the cuID parameter are:
// - name
// - cuname=name
// - WebSphere:cuname=name
// This parameter accepts an incomplete ID as long as the
// incomplete ID can resolve to a unique composition unit
// within the business-level application.
String cuID = "cu1";
cmd.setParameter("cuID", cuID);

System.out.println("\nSet cuID parameter to "
    + cmd.getParameter("cuID"));

// Call the asynchronous client helper to process the command parameters.
try {
    asyncCmdClientHelper.processCommandParameters(cmd);
    System.out.println("\nCompleted process command " +
        "parameters");
} catch (Throwable th) {
    System.out.println("Failed from " +
        "asyncCmdClientHelper.processCommandParameters(cmd).");
    th.printStackTrace();
    System.exit(-1);
}

// Set up the step parameters for the CUOptions step.
// The CUOptions step contains the following arguments that can be edited:
// description - description for the composition unit
// startingWeight - starting weight for the composition unit
//                    within the business-level application.
// startedOnDistributed - to start composition unit upon distribution
//                    to target nodes.
// Valid values are true, false.
// restartBehaviorOnUpdate - restart behavior for the composition
//                    unit when the composition unit is updated.
// Valid values are DEFAULT, ALL, NONE
String stepName = "CUOptions";
CommandStep step = ((TaskCommand) cmd).gotoStep(stepName);

// Composition Unit description:
String description = "cu1 description changed in editCompUnit";

for(int i = 0; i < step.getNumberOfRows(); i++) {
    // Use the following code to change the composition unit step parameters
    // of the CUOptions step. Change your set of step parameters
    // as required by your scenario.

    // For example, set the description.
    step.setParameter("description", description, i);
    System.out.println("\nSet description parameter to " +
        step.getParameter("description", i));
}

// Set up the step parameters for the MapTargets step
stepName = "MapTargets";
step = ((TaskCommand) cmd).gotoStep(stepName);

// In this step the server parameter is required.
// server - target(s) to deploy the composition unit. The default is server1.

```

```

//      To add an additional target to the existing
//      target, add a prefix to the target with a "+". To
//      delete an existing target, add a prefix to the
//      target with a "#". To replace the existing
//      target, use the regular syntax as in the addCompUnit example.
// Example: server = "#server1+server2";
//          String server = "server1";

for(int i = 0; i < step.getNumberOfRows(); i++) {
    // Use the following code to set the server parameter of the MapTargets step.
    // Change your set of step parameters as required by your
    // scenario.

    // For example, set the server.
    step.setParameter("server", server, i);
    System.out.println("\nSet server parameter to " +
        step.getParameter("server", i));
}

// If the RelationshipOptions step is available, the selected
// deployable units of the source asset of the "primary" composition
// unit (that is, the composition unit being added) have dependencies
// on other assets for which there are matching "secondary" composition
// units in the business-level application. The RelationshipOptions step is much like
// CreateAuxCUOptions except that the required secondary composition
// units already exist. Also, each RelationshipOptions row maps one
// deployable unit to one or more secondary composition units, whereas,
// each CreateAuxCUOptions row maps one deployable unit to one
// asset dependency.
//
// Each RelationshipOptions row corresponds to one deployable unit
// with one or more dependency relationships and consists of
// parameter values for the dependency relationships. Some parameters
// are read-only and some of them are editable. To edit parameter
// values, use the same approach as that used to edit parameter values
// in the CUOptions step.
//
// The parameters for this step include:
//
// deplUnit - The name of the deployable unit which has the
//            dependency. (Read-only.)
// relationship - Composition unit dependencies in the form of a
//                list of composition unit IDs. Composition unit
//                IDs are separated by a "plus" sign (+). Each ID
//                can be fully or partially formed as shown with the
//                following examples:
//                WebSphere:cuname=SharedLib1.jar
//                WebSphere:cuname=SharedLib.jar
//                SharedLib.jar
// matchTarget - Specifies whether the server target for the secondary
//               composition units are to match the server target for
//               the primary composition unit. The default value
//               is "true". If the value is set to "false", the
//               secondary composition unit will be created with no
//               target. The target on the secondary composition unit
//               can be set at a later time with the editCompUnit
//               command.
// for(int i = 0; i < step.getNumberOfRows(); i++) {
//     // Use the following if statement to set the relationship and matchTarget parameters
//     // of the RelationshipOptions step. Change your set of
//     // step parameters as required by your scenario.

//     // Uncomment the following code to match the deplUnit and then set
//     // the relationship differently.
//     //String deplUnit = (String) step.getParameter("deplUnit",
//     //                                             i);

```

```

        //if (deplUnit.equals("a1.jar") {
            // For example, change the relationship for the a1.jar file.
            //step.setParameter("relationship", relationship, i);
            //System.out.println("\nSet relationship parameter " +
            //    "to " + step.getParameter("relationship", i));

            // For example, change matchTarget.
            //step.setParameter("matchTarget", matchTarget, i);
            //System.out.println("\nSet matchTarget parameter to "+
            //    step.getParameter("matchTarget", i));
        //}

// The addCompUnit command contains thr ActivationPlanOptions step.
// The user can set the ActivationPlanOptions step parameters similar to
// the step parameters for the CUOptions step in the previous examples.
// The arguments for this step include:
// deplUnit – deployable unit URI (read only parameter)
// activationPlan - specifies a list of runtime components in the
// format of specname=xxxx

// Run the command command to edit the composition unit.
asyncCmdClientHelper.execute(cmd);
System.out.println("\nCompleted running of the command");

// Check the command result.
CommandResult result = cmd.getCommandResult();
if (result != null) {
    if (result.isSuccessful()) {
        System.out.println("\nCommand ran successfully "
            + "with result\n" + result.getResult());
    }
    else {
        System.out.println("\nCommand ran with " +
            "Exception");
        result.getException().printStackTrace();
    }
}
} catch (Exception e) {
    e.printStackTrace();
}
}
}

package com.ibm.ws.management.application.task;

import com.ibm.websphere.management.cmdframework.provider.CommandNotification;
import com.ibm.websphere.management.async.client.AsyncCommandHandlerIF;

public class AsyncCmdTaskHandler implements AsyncCommandHandlerIF {

    public void handleNotification(CommandNotification notification) {
        // Add your own code here to handle the received notification
        System.out.println("\nEXAMPLE: notification received: " +
            notification);
    }
}

```

What to do next

After you edit the composition unit, you can run the updated business-level application.

Deleting a business-level application using programming

You can delete a business-level application using programming. You might delete a business-level application if it is not functioning correctly, it is no longer needed, and so on.

Before you begin

This task assumes a basic familiarity with command framework programming. Read about command framework programming in the application programming interfaces documentation.

Before you can delete a business-level application, you must have created an empty business-level application. You can optionally have added assets or business-level applications as composition units to the empty business-level application. All the composition units in the business-level application must be deleted using the `deleteCompUnit` command before you delete the business-level application. Other business-level applications cannot reference the business-level application that you are deleting. Otherwise, the deletion fails.

You can delete a business-level application using programming, the administrative console, or the `wsadmin` tool.

About this task

You must specify the `blaID` parameter of the business-level application that you are deleting. You might delete a business-level application if it is not functioning correctly, no longer needed, and so on.

Perform the following steps to delete a business-level application using programming.

Procedure

1. Connect to the application server.
The command framework allows the administrative command to be created and run with or without being connected to the application server. This step is optional if the application server is not running.
2. Create the command manager.
The command manager provides the functionality to create a new administrative command or query existing administrative commands.
3. Optionally create the asynchronous command handler for listening to command notifications.
Business-level application commands are implemented as asynchronous commands. To monitor the progress of the running command, you have to create an asynchronous command handler to receive notifications that the command generates.
4. Create the asynchronous command client.
An asynchronous command client provides a higher level interface to work with an asynchronous command. If you created an asynchronous command handler in the previous step, the handler is passed to the asynchronous command client. The asynchronous command client forwards the command notification to the handler and helps to control running of the command.
5. Use the command manager you created in a previous step to create and set up the command that deletes the business-level application.
The command name is `DeleteBLA`. The `blaID` parameter is a required parameter to specify the business-level application to delete.
6. Call the `processCommandParameters` method in the asynchronous command client to process the command parameters.
The command framework asynchronous command model requires this call.
7. Call the asynchronous command client to run the command that deletes the business-level application.

You might have created an asynchronous command handler to implement the `AsyncCommandHandlerIF` interface class in a previous step. If you did, the asynchronous command client listens to command notifications and forwards the notifications to the handler. The handler performs any necessary actions while waiting for the command to complete.

8. Check the command result when the command completes.

When the command finishes running, control is returned to the caller. You can then check the result by calling the `command.getCommandResult` method.

Results

After you successfully run the code, the business-level application is deleted.

Example

The following example shows how to delete a business-level application based on the previous steps. Some statements are split on multiple lines for printing purposes.

```
package com.ibm.ws.management.application.task;

import java.util.Properties;

import com.ibm.websphere.management.AdminClient;
import com.ibm.websphere.management.AdminClientFactory;
import com.ibm.websphere.management.Session;
import com.ibm.websphere.management.cmdframework.AdminCommand;
import com.ibm.websphere.management.cmdframework.CommandMgr;
import com.ibm.websphere.management.cmdframework.CommandResult;
import com.ibm.websphere.management.cmdframework.CommandStep;
import com.ibm.websphere.management.cmdframework.TaskCommand;
import com.ibm.websphere.management.async.client.AsyncCommandClient;

public class DeleteBLA {

    public static void main(String [] args) {

        try {

            // Connect to the application server.
            // This step is optional if you use the local
            // command manager. Comment out the lines to and including
            // CommandMgr cmdMgr = CommandMgr.getClientCommandMgr(
            // soapClient);
            // to get the soapClient soap client if you use the local
            // command manager.

            String host = "localhost";
            String port = "8880"; // Change to your port number if it is
                // not 8880.

            Properties config = new Properties();
            config.put(AdminClient.CONNECTOR_HOST, host);
            config.put(AdminClient.CONNECTOR_PORT, port);
            config.put(AdminClient.CONNECTOR_TYPE,
                AdminClient.CONNECTOR_TYPE_SOAP);
            System.out.println("Config: " + config);
            AdminClient soapClient =
                AdminClientFactory.createAdminClient(config);

            // Create the command manager.
            CommandMgr cmdMgr = CommandMgr.getClientCommandMgr(soapClient);

            // Comment out the previous lines to create a client command
            // manager if you are using a local command manager.
```

```

// Uncomment the following line to create a local command
// manager:
//
// CommandMgr cmdMgr = CommandMgr.getCommandMgr();
System.out.println("\nCreated command manager");

// Optionally create an asynchronous command handler.
// Comment out the following line if no further handling
// of command notification is required:
AsyncCmdTaskHandler listener = new AsyncCmdTaskHandler();

// Create an asynchronous command client.

// Setup the session.
String id = Long.toHexString(System.currentTimeMillis());
String user = "content" + id;
Session session = new Session(user, true);

// If no command handler is used, replace listener with
// null for the following AsyncCommandClient object:
AsyncCommandClient asyncCmdClientHelper = new
AsyncCommandClient(session, listener);
System.out.println("\nCreated async command client");

// Create the command that deletes the business-level application.
String cmdName = "deleteBLA";
AdminCommand cmd = cmdMgr.createCommand(cmdName);
cmd.setConfigSession(session); // Delete the business-level
// application using the session created.
System.out.println("\nCreated " + cmdName);

// Set the blaID parameter to the business-level application to delete.
// Examples of valid formats for the blaID parameter are:
// - bName
// - blaname=bName
// - WebSphere:blaname=bName
// This parameter accepts an incomplete ID as long as the incomplete
// ID can resolve to a unique business-level application.
String blaID = "bla1";
cmd.setParameter("blaID", blaID);

System.out.println("\nSet blaID parameter to "
+ cmd.getParameter("blaID"));

// Call the asynchronous client helper to process parameters.
try {
    asyncCmdClientHelper.processCommandParameters(cmd);
    System.out.println("\nCompleted process command " +
        "parameters");
} catch (Throwable th) {
    System.out.println("Failed from " +
        "asyncCmdClientHelper.processCommandParameters(cmd).");
    th.printStackTrace();
    System.exit(-1);
}

// Call the asynchronous command client to run the command.
asyncCmdClientHelper.execute(cmd);
System.out.println("\nCompleted running of the command");

// Check the command result.
CommandResult result = cmd.getCommandResult();
if (result != null) {
    if (result.isSuccessful()) {
        System.out.println("\nCommand ran successfully "
+ "with result\n" + result.getResult());
    }
}

```

```

        }
        else {
            System.out.println("\nCommand ran with " +
                               "Exception");
            result.getException().printStackTrace();
        }
    }
} catch (Exception e) {
    e.printStackTrace();
}
}
}

package com.ibm.ws.management.application.task;

import com.ibm.websphere.management.cmdframework.provider.CommandNotification;
import com.ibm.websphere.management.async.client.AsyncCommandHandlerIF;

public class AsyncCmdTaskHandler implements AsyncCommandHandlerIF {

    public void handleNotification(CommandNotification notification) {
        // Add your own code here to handle the received notification
        System.out.println("\nEXAMPLE: notification received: " +
                           notification);
    }
}

```

What to do next

You can complete other tasks associated with business-level applications, such as creating other business-level applications, stopping and starting business-level applications, and so on.

Deleting a composition unit using programming

You can delete a composition unit from a business-level application if the composition unit is not functioning correctly, the composition unit is no longer needed, and so on. A composition unit is typically created from a business-level application or an asset and contains configuration information that makes the asset runnable.

Before you begin

This task assumes a basic familiarity with command framework programming. Read about command framework programming in the application programming interfaces documentation.

Before you can delete a composition unit, you must have created an empty business-level application, imported an asset, and added a composition unit to the business-level application. If other composition units depend on the composition unit that you are deleting and you do not use the force option, the deletion fails.

About this task

You can delete a composition unit using programming, the administrative console, or the wsadmin tool. This topic describes how to delete a composition unit using programming.

You must provide the blaID and cuID parameters to specify the composition unit that you are deleting from the business-level application.

Perform the following tasks to delete a composition unit using programming.

Procedure

1. Connect to the application server.
The command framework allows the administrative command to be created and run with or without being connected to the application server. This step is optional if the application server is not running.
2. Create the command manager.
The command manager provides the functionality to create a new administrative command or query existing administrative commands.
3. Optionally create the asynchronous command handler for listening to command notifications.
Business-level application commands are implemented as asynchronous commands. To monitor the progress of the running command, you have to create an asynchronous command handler to receive notifications that the command generates.
4. Create the asynchronous command client.
An asynchronous command client provides a higher level interface to work with an asynchronous command. If you created an asynchronous command handler in the previous step, the handler is passed to the asynchronous command client. The asynchronous command client forwards the command notification to the handler and helps to control running of the command.
5. Use the command manager that you created in a previous step to create and set up the command that deletes a composition unit.
The command name is `deleteCompUnit`. The `blalD` and `culD` parameters are required parameters. The `culD` parameter is used to specify the composition unit to delete from the business-level application, which is specified with the `blalD`. You can optionally provide the `force` parameter to force the deletion if other composition units depend on this composition unit.
6. Call the `processCommandParameters` method in the asynchronous command client to process the command parameters.
The command framework asynchronous command model requires this call.
7. Call the asynchronous command client to run the command that deletes a composition unit.
You might have created an asynchronous command handler to implement the `AsyncCommandHandlerIF` interface class in a previous step. If you did, the asynchronous command client listens to command notifications and forwards the notifications to the handler. The handler performs any necessary actions while waiting for the command to complete.
8. Check the command result when the command completes.
When the command finishes running, control is returned to the caller. You can then check the result by calling the `command.getCommandResult` method.

Results

After you successfully run the code, the composition unit is deleted.

Example

The following example shows how to delete a composition unit from a business-level application based on the previous steps. Some statements are split on multiple lines for printing purposes.

```
package com.ibm.ws.management.application.task;

import java.util.Properties;

import com.ibm.websphere.management.AdminClient;
import com.ibm.websphere.management.AdminClientFactory;
import com.ibm.websphere.management.Session;
import com.ibm.websphere.management.cmdframework.AdminCommand;
import com.ibm.websphere.management.cmdframework.CommandMgr;
import com.ibm.websphere.management.cmdframework.CommandResult;
import com.ibm.websphere.management.cmdframework.CommandStep;
import com.ibm.websphere.management.cmdframework.TaskCommand;
```

```

import com.ibm.websphere.management.async.client.AsyncCommandClient;

public class DeleteCompUnit {

    public static void main(String [] args) {

        try {

            // Connect to the application server.
            // This step is optional if you use the local
            // command manager. Comment out the lines to and including
            // CommandMgr cmdMgr = CommandMgr.getClientCommandMgr(soapClient);
            // to get the soapClient soap client if
            // you use the local command manager.

            String host = "localhost";
            String port = "8880"; // Change to your port number if it is
            // not 8880.

            Properties config = new Properties();
            config.put(AdminClient.CONNECTOR_HOST, host);
            config.put(AdminClient.CONNECTOR_PORT, port);
            config.put(AdminClient.CONNECTOR_TYPE,
                AdminClient.CONNECTOR_TYPE_SOAP);
            System.out.println("Config: " + config);
            AdminClient soapClient =
                AdminClientFactory.createAdminClient(config);

            // Create the command manager.
            CommandMgr cmdMgr = CommandMgr.getClientCommandMgr(soapClient);

            // Comment out the previous lines to create a client command
            // manager if you are using a local command manager.
            // Uncomment the following line to create a local command
            // manager:
            //
            // CommandMgr cmdMgr = CommandMgr.getCommandMgr();
            System.out.println("\nCreated command manager");

            // Optionally create an asynchronous command handler.
            // Comment out the following line if no further handling
            // of command notification is required:
            AsyncCmdTaskHandler listener = new AsyncCmdTaskHandler();

            // Create an asynchronous command client.

            // Set up the session.
            String id = Long.toHexString(System.currentTimeMillis());
            String user = "content" + id;
            Session session = new Session(user, true);

            // If no command handler is used, replace listener with
            // null for the following AsyncCommandClient object:
            AsyncCommandClient asyncCmdClientHelper = new
            AsyncCommandClient(session, listener);
            System.out.println("\nCreated async command client");

            // Create the command that deletes the composition unit.
            String cmdName = "deleteCompUnit";
            AdminCommand cmd = cmdMgr.createCommand(cmdName);
            cmd.setConfigSession(session); // Delete the composition unit from
            // the business-level application
            // using the session created.
            System.out.println("\nCreated " + cmdName);
        }
    }
}

```

```

// Set the blaID parameter to the business-level application with
// the composition unit to delete.
// Examples of valid formats for the blaID parameter are:
// - bName
// - blaname=bName
// - WebSphere:blaname=bName
// This parameter accepts an incomplete ID as long as
// the incomplete ID can resolve to a unique
// business-level application.
String blaID = "bla1";
cmd.setParameter("blaID", blaID);

System.out.println("\nSet blaID parameter to "
    + cmd.getParameter("blaID"));

// Set the cuID parameter to the composition unit that is to be
// deleted.
// Examples of valid formats for the cuID parameter are:
// - name
// - cuname=name
// - WebSphere:cuname=name
// This parameter accepts an incomplete ID as long as the
// incomplete ID can resolve to a unique composition unit
// within the business-level application.
String cuID = "cu1";
cmd.setParameter("cuID", cuID);

System.out.println("\nSet cuID parameter to "
    + cmd.getParameter("cuID"));
// Uncomment the following line of code to set the force parameter
// to force the deletion even if other composition units depend
// on this composition unit.
//
// cmd.setParameter("force", "true");

// Call the asynchronous client helper to process parameters.
try {
    asyncCmdClientHelper.processCommandParameters(cmd);
    System.out.println("\nCompleted process command " +
        "parameters");
} catch (Throwable th) {
    System.out.println("Failed from " +
        "asyncCmdClientHelper.processCommandParameters(cmd).");
    th.printStackTrace();
    System.exit(-1);
}

// Call the asynchronous command client to run the command.
asyncCmdClientHelper.execute(cmd);
System.out.println("\nCompleted running of the command");

// Check the command result.
CommandResult result = cmd.getCommandResult();
if (result != null) {
    if (result.isSuccessful()) {
        System.out.println("\nCommand ran successfully "
            + "with result\n" + result.getResult());
    }
    else {
        System.out.println("\nCommand ran with " +
            "Exception");
        result.getException().printStackTrace();
    }
}
} catch (Exception e) {
    e.printStackTrace();
}

```

```
    }  
  }  
}  
  
package com.ibm.ws.management.application.task;  
  
import com.ibm.websphere.management.cmdframework.provider.CommandNotification;  
import com.ibm.websphere.management.async.client.AsyncCommandHandlerIF;  
  
public class AsyncCmdTaskHandler implements AsyncCommandHandlerIF {  
    public void handleNotification(CommandNotification notification) {  
        // Add your own code here to handle the received notification  
        System.out.println("\nEXAMPLE: notification received: " +  
            notification);  
    }  
}
```

What to do next

You can complete other tasks associated with the business-level application, such as adding or deleting other composition units, listing composition units, and so on.

Chapter 13. Troubleshooting deployment

When you are having problems deploying an application, perform some basic diagnostics and verify your system configuration to solve the problem.

Before you begin

Try to install your application on a product server or cluster. Ensure that your application can be installed to the deployment target. For example, if your application contains modules that support Java Platform, Enterprise Edition (Java EE) 6 or use a Version 8 product feature or API, you must install the application to a Version 8 deployment target.

About this task

Determine which of the following steps apply to the deployment problem and read the suggested topics.

Procedure

- If you cannot install the application, troubleshoot problems deploying applications.
See the topics on application deployment problems and troubleshooting tips.
- If you can install the application but it does not start, troubleshoot problems starting applications.
See the topics on application deployment and startup problems.
- If your application contains many classes with annotations and takes a long time to deploy, reduce annotation searches to speed up deployment.
See the topic on reducing annotation searches during application deployment.
- If you cannot uninstall the application, see the topic on application uninstallation problems.

What to do next

If the topics in this information center do not resolve the deployment problem, examine current information available from IBM Support on known problems and their resolution. IBM Support has documents that can save you time gathering information needed to resolve this problem. Before opening a PMR, see topic on troubleshooting help from IBM.

Application deployment problems

You might encounter problems when deploying, installing, or promoting applications. This topic suggests ways to resolve the problems.

What kind of problem are you having?

- “I installed my application using the wsadmin tool, but the application does not display under Applications > Application Types > WebSphere enterprise applications” on page 558
- “Unable to save a deployed application” on page 558
- “WASX7015E error running wsadmin command \$AdminApp installInteractive or \$AdminApp install” on page 559
- “Cannot install a CMP or BMP entity bean in an EJB 3.0 module” on page 559
- “Data definition language (DDL) generated by an assembly tool throws SQL error on target platform” on page 559
- “Error message ADMA0004E: Validation error in task Specifying the Default Datasource for EJB Modules returned when installing application using the administrative console or the wsadmin tool” on page 560
- “Cannot load resource WEB-INF/ibm-web-bnd.xmi in archive file” on page 560

- “Error message No valid target is specified in ObjectName anObject for module module_name from installation ” on page 561
- “addNode -includeapps option does not appear to upload all applications to the deployment manager” on page 561
- ““Timeout!!!” error displays when attempting to install an enterprise application in the administrative console ” on page 561
- “I get a NameNotFoundException message when deploying an application that contains an EJB module” on page 561
- “I get compilation errors and EJB deploy fails when installing an EJB JAR file generated for Version 5.x or earlier” on page 562
- “While uploading documents, addNode -includeapps fails with an OutOfMemoryError exception” on page 562
- “I get an OutOfMemory exception in the deployment manager” on page 562
- “After installing the application onto a different machine, the application does not run” on page 562
- “A single file replaces all application files during application update” on page 563

Check the following first:

- Verify that the logical name that you have specified to appear on the console for your application, enterprise bean module or other resource does not contain invalid characters such as these: - / \ : * ? " < > |.
- If the application was installed using the wsadmin \$AdminApp install command with the **-local** flag, restart the server or rerun the command without the `-local` flag.

If you do not see a problem that resembles yours, or if the information provided does not solve your problem, check to see if the problem is identified and documented.

I installed my application using the wsadmin tool, but the application does not display under Applications > Application Types > WebSphere enterprise applications

The application might be installed but you have not saved the configuration:

1. Verify that the application subdirectory is located under the *app_server_root/installedApps* directory.
2. Run the \$AdminApp list command and verify that the application is not among those displayed.
 - In the bin directory, run the wsadmin.bat or wsadmin.sh command.
 - From the wsadmin prompt, enter \$AdminApp list and verify that the problem application is not among the items that display.
3. Reinstall your application using the wsadmin tool. Run the \$AdminConfig save command in the wsadmin tool before exiting.

Unable to save a deployed application

If you are unable to save a deployed application, the problem might be that too many files are opened, exceeding the limit of the operating system.

Only root has authority to adjust the maximum number of files for each process. Complete the following steps to modify the application to close files with disciplines:

1. After you open a file and complete your work, call the close method of the file to release the file handle back to the operating system.
2. Using the java.io.FileInputStream and the FileOutputStream classes as examples, you can invoke their close method to release any system resources that are associated with the stream.

WASX7015E error running wsadmin command \$AdminApp installInteractive or \$AdminApp install

This problem has two possible causes:

- If the full text of the error is similar to:

```
WASX7015E: Exception running command:
"$AdminApp installInteractive Documents and Settings/
myUserName/Desktop/MyApp/myapp.ear"; exception information:
com.ibm.bsf.BSFException: error while
evaluating Jacl expression: can't find method "installInteractive"
with 3 argument(s) for class
"com.ibm.ws.scripting.AdminAppClient"
```

The file and path name are incorrectly specified. In this case, since the path included spaces, it was interpreted as multiple parameters by the wsadmin program.

Enter the path of the .ear file correctly. In this case, by enclosing it in double quotes:

```
$AdminApp installInteractive
"Documents and Settings/myUserName/Desktop/MyApps/myapp.ear"
```

- If the full text of the error is similar to:

```
WASX7015E: Exception running command: "$AdminApp installInteractive
MyApps\myapp.ear ";
exception information: com.ibm.ws.scripting.ScriptingException: WASX7115E:
Cannot read input file
"WebSphere\AppServer\bin\MyAppsmyapp.ear"
```

The application path is incorrectly specified. In this case, you must use "forward-slash" (/) separators in the path.

Cannot install a CMP or BMP entity bean in an EJB 3.0 module

When installing an EJB 3.0 module that contains a container-managed persistence (CMP) or bean-managed persistence (BMP) entity bean, the installation fails.

The product does not support installation of applications that have a CMP or BMP entity bean packaged in an EJB 3.0 module. You must package CMP or BMP entity beans in an EJB 2.1 or earlier module.

To resolve this problem:

1. Package the CMP or BMP entity beans in EJB 2.1 or earlier modules.
2. Try installing your application with the EJB 2.1 or earlier modules.

Data definition language (DDL) generated by an assembly tool throws SQL error on target platform

If you receive SQL errors in attempting to execute data definition language (DDL) statements generated by an assembly tool on a different platform, for example if you are deploying a container-managed persistence (CMP) enterprise bean designed on Windows onto a UNIX operating system server, try the following actions:

- Browse the DDL statements for dependencies on specific user identifiers and passwords, and correct as necessary.
- Browse the DDL statements for dependencies on specific server names, and correct as necessary.
- Refer to the message reference of the vendor for causes and suggested actions regarding specific SQL errors. For IBM DB2, you can view the message references online at <http://www.ibm.com/cgi-bin/db2www/data/db2/udb/winos2unix/support/index.d2w/report>.

If you receive the following error after executing a DDL file created on the Windows operating system or on operating systems such as AIX® or Linux, the problem might come from a difference in file formats:

```
SQL0104N  An unexpected token "CREATE TABLE AGENT (COMM DOUBLE, PERCENT DOUBLE, P"
was found following "      ".  Expected tokens may include: " ".
SQLSTATE=42601
```

To resolve this problem:

- Edit the DDL in the vi editor, removing the Ctl-M character at the beginning of each line.

Error message ADMA0004E: Validation error in task Specifying the Default Datasource for EJB Modules returned when installing application using the administrative console or the wsadmin tool

If you see the following error when trying to install an application through the administrative console or the wsadmin command prompt:

```
AppDeploymentException: [ADMA0014E: Validation failed.
ADMA0004E: Validation error in task Specifying the Default Datasource for
EJB Modules  JNDI name is not
specified for module beanameBean Jar with URI filename.jar,META-INF/ejb-jar.xml.
You have not specified the
data source for each CMP bean belonging to this module. Either specify the data
source for each CMP beans or
specify the default data source for the entire module.]
```

one possible cause is that, in WebSphere Application Server Version 4.0, it was mandatory to have a data source defined for each CMP bean in each JAR. In Version 5.0 and later releases, you can specify either a data source for a container-managed persistence (CMP) bean or a default data source for all CMP beans in the JAR file. Thus during installation interaction, such as the installation wizard in the administrative console, the data source fields are optional, but the validation performed at the end of the installation checks to see that at least one data source is specified.

To correct this problem, step through the installation again, and specify either a default data source or a data source for each CMP-type enterprise bean.

If you are using the wsadmin tool, use the \$AdminApp installInteractive *filename* command to receive prompts for data sources during installation, or to provide them in a response file.

Cannot load resource WEB-INF/ibm-web-bnd.xmi in archive file

The web application tmp.war installs on WebSphere Application Server Versions 5.0 and 5.1, but fails on a WebSphere Application Server Version 6.0 or later server. The application fails to install because the WEB-INF/ibm-web-bnd.xmi file contains xmi tags that the underlying WCCM model no longer recognizes.

The following error messages display:

```
IWAE0007E Could not load resource "WEB-INF/ibm-web-bnd.xmi" in archive "tmp.war"
[2/24/05 14:53:10:297 CST] 000000bc SystemErr      R
AppDeploymentException:
com.ibm.etools.j2ee.commonarchivecore.exception.ResourceLoadException:
IWAE0007E Could not load resource "WEB-INF/ibm-web-bnd.xmi" in archive "tmp.war"
[2/24/05 14:53:10:297 CST] 000000bc SystemErr      R
com.ibm.etools.j2ee.commonarchivecore.exception.ResourceLoadException:
IWAE0007E Could not load resource "WEB-INF/ibm-web-bnd.xmi" in archive "tmp.war"
!Stack_trace_of_nested_exce!
com.ibm.etools.j2ee.exception.WrappedRuntimeException: Exception occurred loading
WEB-INF/ibm-web-bnd.xmi
!Stack_trace_of_nested_exce!
```

To work around this problem, remove the xmi:type=EJBLocalRef tag from the ibm-web-bnd.xmi file. Removing this tag does not affect the application because the tag was previously used for matching the cross document reference type. The application now works for the WebSphere Application Server Version 5.1 and later releases.

Note: For IBM extension and binding files, the .xmi or .xml file name extension is different depending on whether you are using a pre-Java EE 5 application or module or a Java EE 5 or later application or module. An IBM extension or binding file is named `ibm-*-ext.xmi` or `ibm-*-bnd.xmi` where `*` is the type of extension or binding file such as `app`, `application`, `ejb-jar`, or `web`. The following conditions apply:

- For an application or module that uses a Java EE version prior to version 5, the file extension must be `.xmi`.
- For an application or module that uses Java EE 5 or later, the file extension must be `.xml`. If `.xmi` files are included with the application or module, the product ignores the `.xmi` files.

However, a Java EE 5 or later module can exist within an application that includes pre-Java EE 5 files and uses the `.xmi` file name extension.

The `ibm-webservices-ext.xmi`, `ibm-webservices-bnd.xmi`, `ibm-webservicesclient-bnd.xmi`, `ibm-webservicesclient-ext.xmi`, and `ibm-portlet-ext.xmi` files continue to use the `.xmi` file extensions.

Error message No valid target is specified in ObjectName *anObject* for module *module_name* from installation

This error can occur in a clustered environment if the target cell, node, server or cluster into which the application is to be installed is incorrectly specified. For example, it can occur if the target is misspelled.

To correct this problem, check the target names against the actual WebSphere Application Server topology and reenter them with corrections.

addNode -includeapps option does not appear to upload all applications to the deployment manager

This error can occur when some or all applications on the target node are already uploaded to the deployment manager. The `addNode` program detects which applications are already installed and does not upload them again.

Use the administrative console to browse the deployment manager configuration and see the applications that are already installed.

"Timeout!!!" error displays when attempting to install an enterprise application in the administrative console

This error can occur if you attempt to install an enterprise application that has not been deployed.

To correct this problem:

- Open the `file_name.ear` file in an assembly tool and then click **Deploy**. This action creates a file with a name like `Deployed_file_name.ear`.
- In the administrative console, install the deployed EAR file.

I get a NameNotFoundException message when deploying an application that contains an EJB module

If you specify that the EJB deployment tool be run during application installation and the installation fails with a `NameNotFoundException` message, ensure that the input JAR or EAR file does not contain source files. If there are source files in the input JAR or EAR file, the EJB deployment tools runs a rebuild before generating the deployment code.

To work around this problem, either remove the source files or include all dependent classes and resource files on the class path. Otherwise, the source files or the lack of access to dependent classes and resource files might cause problems during rebuilding of your application on the server.

I get compilation errors and EJB deploy fails when installing an EJB JAR file generated for Version 5.x or earlier

When installing an old application that uses EJB modules that were built to run on WebSphere Application Server Version 5.x or earlier, compilation errors result and EJB deploy fails. The EJB JAR file contains Java source for the old generated code. The old Java source was generated for Version 5.x or before but, when deployed to a WebSphere Application Server Version 6.x or later product, it is compiled using the Version 6.0 or later runtime JAR files.

To work around this problem, remove all .java files from the application EAR file. After the Java source files are removed, you can deploy the application onto a server successfully.

While uploading documents, addNode -includeapps fails with an OutOfMemoryError exception

This error can occur when you use addNode -includeapps while you are installing applications with large EAR files. To correct this problem:

- If you are using addNode to add a node from the base server, modify the addNode script to include the following parameter:
`-Xmxsize`
- If you are adding a node from the administrative console, increase the *maximumHeapSize* in the Java virtual machine settings of the deployment manager, then restart the deployment manager.

I get an OutOfMemory exception in the deployment manager

If you receive an OutOfMemory exception when trying to install applications with large EAR files, try increasing the maximum heap size of the deployment manager.

Check the options you specified on the Java virtual machine page of the administrative console. Increase the *maximumHeapSize* in the Java virtual machine settings of the deployment manager. Then, restart the deployment manager, and try installing the application again.

If this change does not correct the problem, see "Diagnosing out-of-memory errors and Java Heap memory leaks."

After installing the application onto a different machine, the application does not run

If your application uses application level resources, its application level node information must be correct for the application to run as expected.

When you add application level resources to an application and deploy the application onto a machine, ensure that the application level node information is correct. Otherwise, when you install the application onto a different machine, it is installed to the wrong location and the application does not run as expected.

You can update the application level node information using an assembly tool. Update the nodeName from deploymentTargets of the deployment.xml file under ibmconfig. Also, ensure that binariesURL from deployedObject of the deployment.xml file has the correct path.

A single file replaces all application files during application update

If you select the **Replace or add a single file** option of the application update wizard and the currently deployed application consists of several files, specify the full path name of the file to be replaced or added for **Specify the path beginning with the installed application archive file to the file to be replaced or added**.

A full path name usually has the structure *directory_path/file_name* and resembles the following:

```
PriceChangeSession.jar/priceChangeSession/priceChangeSessionBean.class
```

Do not specify less than the full path name for **Specify the path beginning with the installed application archive file to the file to be replaced or added**. For example, do not specify only a directory path:

```
PriceChangeSession.jar/priceChangeSession
```

If you specify less than a full path name, all files in the directory of the currently deployed application might be replaced by the single new file that was specified under **Specify the path to the file**.

Application deployment troubleshooting tips

When you first test or run a deployed application, you might encounter problems.

Select the problem you are having with testing or the first run of deployed code for WebSphere Application Server:

- “Web resource is not displayed” on page 572.

You can use the following administrative console pages to inspect the configuration of your applications and JMS resources:

- For a view of the JMS resources for a given application, see the following page: `../ae/AppToSIBRefs_DetailForm.dita`.
- For a view of the applications and JMS resources for a given default messaging provider destination, see the following page: `../ae/AppsFromSIBRefs_DetailForm.dita`.

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

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

Application startup errors

Use this information for troubleshooting problems that occur when starting an application.

What kind of error do you see when you start an application?

- “HTTP server and Application Server are working separately, but requests are not passing from HTTP server to Application Server” on page 564
- “File serving problems” on page 565
- “Graphics do not appear in the JSP file or servlet output” on page 565
- “SRVE0026E: [Servlet Error]-[Unable to compile class for JSP file” on page 566
- “Message like “Message: /jspname.jsp(9,0) Include: Mandatory attribute page missing” appears when attempting to browse JSP file” on page 567
- “The JSP Batch Compiler fails with the message “Enterprise Application [application name you typed in] not found.”” on page 567
- “There is a translation problem with non-English browser input” on page 567

- “Scroll bars do not appear around items in the browser window” on page 567
- “Error “Page cannot be displayed... server not found or DNS error” appears when attempting to browse a JavaServer Pages (JSP) file using Internet Explorer” on page 567
- “DeploymentDescriptorLoadException results when starting an application” on page 567

The following note applies to the `ibm-web-ext.xml` references throughout this topic:

Note: For IBM extension and binding files, the `.xml` or `.xmi` file name extension is different depending on whether you are using a pre-Java EE 5 application or module or a Java EE 5 or later application or module. An IBM extension or binding file is named `ibm-*-ext.xml` or `ibm-*-bnd.xml` where `*` is the type of extension or binding file such as `app`, `application`, `ejb-jar`, or `web`. The following conditions apply:

- For an application or module that uses a Java EE version prior to version 5, the file extension must be `.xmi`.
- For an application or module that uses Java EE 5 or later, the file extension must be `.xml`. If `.xmi` files are included with the application or module, the product ignores the `.xmi` files.

However, a Java EE 5 or later module can exist within an application that includes pre-Java EE 5 files and uses the `.xmi` file name extension.

The `ibm-webservices-ext.xml`, `ibm-webservices-bnd.xml`, `ibm-webservicesclient-bnd.xml`, `ibm-webservicesclient-ext.xml`, and `ibm-portlet-ext.xml` files continue to use the `.xmi` file extensions.

HTTP server and Application Server are working separately, but requests are not passing from HTTP server to Application Server

If your HTTP server appears to be functioning correctly, and the Application Server also works on its own, but browser requests sent to the HTTP server for pages are not being served, a problem exists in the WebSphere Application Server plug-in.

In this case:

1. Determine whether the HTTP server is attempting to serve the requested resource itself, rather than forwarding it to the WebSphere Application Server.
 - a. Browse the HTTP server access log (*IHS install root/logs/access.log* for IBM HTTP Server). It might indicate that it could not find the file in its own document root directory.
 - b. Browse the plug-in log file as described below.
2. Refresh the `plugin-cfg.xml` file that determines which requests sent to the HTTP server are forwarded to the WebSphere Application Server, and to which Application Server.

Use the console to refresh this file:

- In the WebSphere Application Server administrative console, expand the Environment tree control.
 - Click **Update WebSphere Plugin**.
 - Stop and restart the HTTP server.
 - Retry the web request.
3. Browse the `plugin_install_root/logs/web_server_name/http_plugin.log` file for clues to the problem. Make sure the timestamps with the most recent plug-in information stanza, which is printed out when the plug-in is loaded, correspond to the time the web server started.
 4. Turn on plug-in tracing by setting the `LogLevel` attribute in the `plugin-cfg.xml` file to `Trace` and reloading the request. Browse the `plugin_install_root/logs/Web_server_name/http_plugin.log` file. You should be able to see the plug-in attempting to match the request URI with the various URI definitions for the routes in the `plugin-cfg.xml`. Check which rules the plug-in is not matching against and then figure out if you need to add additional ones. If you just recently installed the application you might need to manually regenerate the plug-in configuration to pick up the new URIs related to the new application.

For further details on troubleshooting plug-in-related problems, see Webserver plug-in troubleshooting tips located in the *Administering applications and their environment* PDF book.

File serving problems

If text output appears on your JSP- or servlet-supported web page, but image files do not:

- Verify that your files are in the right place: the **document root** directory of your web application WebSphere Application Server follows the J2EE standard, which means that the document root is the *web_module_name.war* directory of your deployed web application.
Typically this directory will be found in the *install_root/installedApps/nodename/appname.ear* directory or *install_root/installedApps/nodename/appnameNetwork.ear* directory.
If the files are in a subdirectory of the document root, verify that the reference to the file reflects that. That is, if the *invoices.html* file is stored in Windows directory *web_module_name.war\invoices*, then links from other pages in the web application to display it should read "*invoices\invoices.html*", not "*invoices.html*".
- Verify that your web application is configured to enable file serving (in other words, that it is enabled to display static resources like image and .html files):
 1. View the file serving property of the hosting web module by browsing the source .war file in an assembly tool. If necessary, update the property and redeploy the module. For more information about the assembly tool, refer to the assembly tools section of the *Developing and deploying applications* PDF book.
 2. Edit the `fileServingEnabled` property in the deployed web application *ibm-web-ext.xmi* configuration file.
The file typically is found in the *install_root/config/cells/nodename* or *nodenameNetwork/applications/application_name/deployments/application_name/Web_module_name/web-inf* directory.

Graphics do not appear in the JSP file or servlet output

If text output appears on your JSP- or -servlet-supported web page, but image files do not:

- Verify that your graphic files are in the right place: the **document root** directory of your web application. The product follows the J2EE standard, which means that the document root is the *web_module_name.war* directory of your deployed web application.
Typically, this directory is found in the *install_root/installedApps/nodename/appname.ear* directory or *install_root/installedApps/nodename/appnameNetwork.ear* directory.
If the graphics files are in a subdirectory of the document root, verify that the reference to the graphic reflects that; for example, if the *banner.gif* file is stored in Windows directory *web_module_name.war/images*, the tag to display it should read: ``, not ``.
- Verify that your web application is configured to enable file serving (that is, display of static resources like image and .html files).
 1. View the file serving property of the hosting web module by browsing the source .war file in an assembly tool. If necessary, update the property and redeploy the module. For more information about the assembly tool, refer to the assembly tools section of the *Developing and deploying applications* PDF book.
 2. Edit the `fileServingEnabled` property in the deployed web application *ibm-web-ext.xmi* configuration file.
The file typically is found in the *install_root/config/cells/nodename* or *nodenameNetwork/applications/application_name/deployments/application_name/Web_module_name/web-inf* directory.
 3. After completing the previous step:
 - In the administrative console, expand the Environment tree control .
 - Click **Update WebSphere Plugin**.
 - Stop and restart the HTTP server and retry the web request.

SRVE0026E: [Servlet Error]-[Unable to compile class for JSP file]

If this error appears in a browser when trying to access a new or modified .jsp file for the first time, the most likely cause is that the JSP file Java source failed (was incorrect) during the javac compilation phase.

Check the log files for a compiler error message.

Note: This topic references one or more of the application server log files. Beginning in WebSphere Application Server Version 8.0 you can configure the server to use the High Performance Extensible Logging (HPEL) log and trace infrastructure instead of using SystemOut.log, SystemErr.log, trace.log, and activity.log files or native z/OS logging facilities. If you are using HPEL, you can access all of your log and trace information using the LogViewer command-line tool from your server profile bin directory. See the information about using HPEL to troubleshoot applications for more information on using HPEL.

Fix the problem in the JSP source file, save the source and request the JSP file again.

If this error occurs when trying to serve a JSP file that was copied from another system where it ran successfully, then there is something different about the new server environment that prevents the JSP file from running. Browse the text of the error for a statement like:

```
Undefined variable or class name: MyClass
```

This error indicates that a supporting class or jar file is not copied to the target server, or is not on the class path. Find the MyClass.class file, and place it on the web module WEB-INF/classes directory, or place its containing .jar file in the Web module WEB-INF/lib directory.

Verify that the URL used to access the resource is correct by doing the following:

- For a JSP file, html file, or image file: `http://host_name/Web_module_context_root/subdir under doc root, if any/filename.ext`. The document root for a web application is the `application_name.WAR` directory of the installed application.
 - For example, to access the myJsp.jsp file, located in `c:\WebSphere\ApplicationServer\installedApps\myEntApp.ear\myWebApp.war\invoices` on myhost.mydomain.com, and assuming the context root for the myWebApp web module is myApp, the URL is `http://myhost.mydomain.com/myApp/invoices/myJsp.jsp`.
 - JSP serving is enabled by default. File serving for HTML and image files must be enabled as a property of the web module, in an assembly tool, or by setting the `fileServingEnabled` property to **true** in the `ibm-web-ext.xmi` file of the installed web application and restarting the application. For more information about the assembly tool, refer to the assembly tools section of the *Developing and deploying applications* PDF book.
- For servlets served by class name, the URL is `http://hostname/Web_module_context_root/servlet/packageName.className`.
- Serving servlets by class name must be enabled as a property of the web module, and is enabled by default. File serving for HTML and image files must be enabled as a property of the Web application, in an assembly tool, or by setting the `fileServingEnabled` property to **true** in the `ibm-web-ext.xmi` file of the installed web application and restarting the application. For more information about the assembly tool, refer to the assembly tools section of the *Developing and deploying applications* PDF book.

Correct the URL in the "from" HTML file, servlet or JSP file. An HREF with no leading slash (/) inherits the calling resource context. For example:

- an HREF in `http://[hostname]/myapp/servlet/MyServlet` to "ServletB" resolves to `"http://hostname/myapp/servlet/ServletB"`
- an HREF in `http://[hostname]/myapp/servlet/MyServlet` to "servlet/ServletB" resolves to `"http://hostname/myapp/servlet/servlet/ServletB"` (an error)
- an HREF in `http://[hostname]/myapp/servlet/MyServlet` to "/ServletB" resolves to `"http://hostname/ServletB"` (an error, if ServletB requires the same context root as MyServlet)

Message like "Message: /jspname.jsp(9,0) Include: Mandatory attribute page missing" appears when attempting to browse JSP file

It is probable that the JSP file failed during the translation to Java phase. Specifically, a JSP directive, in this case an Include statement, was incorrect or referred to a file that could not be found.

To correct this problem, fix the problem in the JSP source, save the source and request the JSP file again.

The JSP Batch Compiler fails with the message "Enterprise Application [application name you typed in] not found."

It is probable that the full enterprise application path and name, starting with the .ear subdirectory that resides in the applications directory is expected as an argument to the JspBatchCompiler tool, not just the display name.

For example:

- "JspBatchCompiler -enterpriseapp.name sampleApp.ear/deployments/sampleApp" is correct, as opposed to
- "JspBatchCompiler -enterpriseapp.name sampleApp", which is incorrect.

There is a translation problem with non-English browser input

If non-English-character-set browser input cannot be translated after being read by a servlet or JSP file, ensure that the request parameters are encoded according to the expected character set before reading. For example, if the site is Chinese, the target .jsp file should have a line:

```
req.setCharacterEncoding("gb2312");
```

before any req.getParameter method calls.

This problem affects servlets and jsp files ported from earlier versions of WebSphere Application Server, which converted characters automatically based upon the locale of the WebSphere Application Server.

Scroll bars do not appear around items in the browser window

In some browsers, tree or list type items that extend beyond their allotted windows do not have scroll bars to permit viewing of the entire list.

To correct this problem, right-click on the browser window and click **Reload** from the menu.

Error "Page cannot be displayed... server not found or DNS error" appears when attempting to browse a JavaServer Pages (JSP) file using Internet Explorer

This error can occur when an HTTP timeout causes the servant to be brought down and restarted. To correct this problem, increase the ConnectionIOTimeout value:

1. From the administrative console, select **System administration > Deployment manager > Administration Services > Custom Properties**
2. Select ConnectionIOTimeout.
3. Increase the ConnectionIOTimeout value.
4. Click **OK**.

DeploymentDescriptorLoadException results when starting an application

If an application server on which the application is deployed synchronizes configuration with the deployment manager during server startup, then the application might not start and a

DeploymentDescriptorLoadException might be written to the server SystemErr.log file. Stop and restart the server, and then try starting the application again.

Application startup problems

When an application is not starting or starting with errors, the problem could be from one of various sources.

What kind of error do you see when you start an application?

- “WSVR0100W: An error occurred initializing, application_name java.lang.NullPointerException when starting a migrated application”
- A “java.lang.ClassNotFoundException: classname Bean_AdderServiceHome_04f0e027Bean” on page 569 error occurs
- A “ConnectionFactory E J2CA0102E: Invalid EJB component: Cannot use an EJB module with version 1.1 using The Relational Resource Adapter” on page 569 error occurs
- “NMSV0605E: "A Reference object looked up from the context..." error when starting an application” on page 570.
- “A Page Not Found, Array Index Out of Bounds, or other error when an updated application restarts” on page 570

If none of these errors match the error you see:

- Browse the log files of the application server for this application looking for clues. By default, these files are: *app_server_root/logs/server_name/SystemErr.log* and *SystemOut.log*.
- Look up any error or warning messages in the message reference table by clicking the **Reference** view and expanding **Messages**.

Note: This topic references one or more of the application server log files. Beginning in WebSphere Application Server Version 8.0 you can configure the server to use the High Performance Extensible Logging (HPEL) log and trace infrastructure instead of using *SystemOut.log*, *SystemErr.log*, *trace.log*, and *activity.log* files or native z/OS logging facilities. If you are using HPEL, you can access all of your log and trace information using the LogViewer command-line tool from your server profile bin directory. See the information about using HPEL to troubleshoot applications for more information on using HPEL.

If you do not see a problem that resembles yours, or if the information provided does not solve your problem, see IBM Support troubleshooting information.

WSVR0100W: An error occurred initializing, *application_name* java.lang.NullPointerException when starting a migrated application

After you migrate an enterprise application to Version 8.0, the application might not start. Attempts to start the application result in an error such as WSVR0100W: An error occurred initializing, *application_name* java.lang.NullPointerException.

Examine the *deployment.xml* file of the migrated application, and remove *targetMapping* statements such as the following:

```
<targetMappings xmi:id="DeploymentTargetMapping_1279594183813" enable="true"/>
```

Then, try starting the application again. The Version 8.0 runtime has an application validation process that might not support migrated *targetMappings* settings.

java.lang.ClassNotFoundException: *classname* Bean_AdderServiceHome_04f0e027Bean

An similar exception occurs when you try to start an undeployed application containing enterprise beans, or containing undeployed enterprise bean modules.

To avoid this problem, do the following:

- Use an assembly tool and administrative console to generate deployment code and install the application or Enterprise JavaBeans module onto a server.
 1. Uninstall the application or Enterprise JavaBeans module in the administrative console.
 2. Configure your assembly tool so the target server is a WebSphere Application Server installation. If you do not have access to the target server, you can specify a false location such as /temp. Specifying a false location enables you to assemble and generate deployment code for the enterprise bean.
 3. In the Project Explorer view of an assembly tool, right-click the enterprise bean (Enterprise JavaBeans) in the undeployed .ear file containing the Enterprise JavaBeans module or the stand-alone undeployed Enterprise JavaBeans JAR file, and click **Deploy**. If your assembly tool can access the WebSphere Application Server target server, deployment code is generated for the Enterprise JavaBeans and the assembly tool attempts to install the application or module onto the target server. If your assembly tool cannot access the WebSphere Application Server target server or the installation fails, use the deployment code that is generated for the next step.
For information on using an assembly tool, refer to the topic on assembling applications.
 4. Use the `wsadmin $AdminApp install` command or the administrative console to install the deployed version created by the assembly tool.
- If you use the `wsadmin $AdminApp install` command, uninstall it and then reinstall using the `-EJBDeploy` option. Follow the install command with the `$AdminConfig save` command.

ConnectionFactory J2CA0102E: Invalid EJB component: Cannot use an EJB module with version 1.1 using The Relational Resource Adapter

This error occurs when an enterprise bean developed to the Enterprise JavaBeans 1.1 specification is deployed with a WebSphere Application Server V5 J2C-compliant data source, which is the default data source. By default, persistent enterprise beans created under WebSphere Application Server V4.0 using the Application Assembly Tool fulfill the Enterprise JavaBeans 1.1 specification. To run on WebSphere Application Server V6, these enterprise beans must be associated with a WebSphere Application Server V4.0-type data source.

Either modify the mapping in the application of enterprise beans to associate 1.x container managed persistence (CMP) beans to associate them with a V4.0 data source or delete the existing data source and create a V4.0 data source with the same name.

To modify the mapping in the application of enterprise beans, in the WebSphere Application Server administrative console, select the properties for the problem application and use **Map resource references to resources** or **Map data sources for all 1.x CMP beans** to switch the data source the enterprise bean uses. Save the configuration and restart the application.

To delete the existing data source and create a V4.0 data source with the same name:

1. In the administrative console, click **Resources > Manage JDBC Providers > JDBC_provider_name > Data sources**.
2. Delete the data source associated with the Enterprise JavaBeans 1.1 module.
3. Click **Resources > Manage JDBC Providers > JDBC_provider_name > Data sources (Version 4)**.
4. Create the data source for the Enterprise JavaBeans 1.1 module.
5. Save the configuration and restart the application.

NMSV0605E: "A Reference object looked up from the context..." error when starting an application

If the full text of the error is similar to:

```
[7/17/02 15:20:52:093 CDT] 5ae5a5e2 UrlContextHel W NMSV0605E:
A Reference object looked up from the context
"java:" with the name "comp/PM/WebSphereCMPConnectionFactory" was sent to the JNDI Naming Manager
and an exception resulted. Reference data follows:
Reference Factory Class Name: com.ibm.ws.naming.util.IndirectJndiLookupObjectFactory
Reference Factory Class Location URLs:
Reference Class Name: java.lang.Object
Type: JndiLookupInfo
Content: JndiLookupInfo: ; jndiName="eis/jdbc/MyDatasource_CMP"; providerURL="";
initialContextFactory=""
```

then the problem might be that the data source intended to support a CMP enterprise bean is not correctly associated with the enterprise bean.

To resolve this problem:

1. Select the **Use this Data Source in container managed persistence (CMP)** check box in the data source "General Properties" panel of the administrative console.
2. Verify the JNDI name in either of the following ways:
 - Verify that the JNDI name given in the administrative console under **Resources > Manage JDBC Providers > DataSource > JNDI Name** for DataSource matches the JNDI name given for CMP or BMP resource bindings at the time of assembling the application in an assembly tool.
 - Check the JNDI name for CMP or BMP resource bindings specified in the code by J2EE application developer. Open the deployed .ear folder in an assembly tool, and look for the JNDI name for your entity beans under CMP or BMP resource bindings. Verify that the names match.

A Page Not Found, Array Index Out of Bounds, or other error when an updated application restarts

If an application is updated while it is running, WebSphere Application Server automatically stops the application or only its changed components, updates the application logic, and restarts the stopped application or its components. For more information on the restarting of updated applications, refer to Fine-grained recycle behavior in *IBM WebSphere Developer Technical Journal: System management for WebSphere Application Server V6 -- Part 5 Flexible options for updating deployed applications*.

A Page Not Found, Array Index Out of Bounds, or other error might occur during restarting.

To minimize the occurrence of such errors, update applications in a test environment before updating the applications in a production environment. Do not put changes directly into a production environment.

Reducing annotation searches during application deployment

Enterprise applications that contain many classes with annotations might take a long time to deploy. Java EE 5 introduced annotations to add metadata to Java classes. Because of performance issues associated with reflection and because classes are not always loadable at deployment, bytecode scanning technology is used to retrieve annotation metadata. Java EE 5 or later applications with many classes might experience long deployment times because every class within the application is inspected during deployment. You can reduce the number of annotations to inspect by specifying the modules and Java packages to ignore for annotations processing in the `amm.filter.properties` file or by configuring system properties.

Before you begin

Install an application that supports Java Platform, Enterprise Edition (Java EE) 5 or later on a product server or cluster. If deployment is unreasonably slow and you will be deploying this application again in the future, complete a procedure in this topic to reduce the number of classes that are searched for annotations during deployment.

About this task

The product provides a configurable filtering function to reduce the number of classes that are searched for annotations. You can identify which modules or Java packages to ignore for annotations processing through two properties:

- Ignore-Scanning-Archives
- Ignore-Scanning-Packages

A default set of values is provided in the `amm.filter.properties` file in `app_server_root/properties`. The property values provide both coarse and fine grained control over the search scope for annotations processing. Use of the Ignore-Scanning-Archives property reduces deployment time more than use of the Ignore-Scanning-Packages property. The syntax for the Ignore-Scanning-Archives and Ignore-Scanning-Packages properties follows the comma-separated value convention. No wildcard or regular expressions are permitted and values are case-sensitive.

The default set of values can be changed by an administrator or augmented by a user using one of the following steps.

Procedure

- Place an `amm.filter.properties` file in the `profile_root/properties` directory.
 - Use system properties to supply values for the Ignore-Scanning-Archives and Ignore-Scanning-Packages properties.
 - The `com.ibm.ws.amm.scan.context.filter.archives` system property supplies values for the Ignore-Scanning-Archives property.
 - The `com.ibm.ws.amm.scan.context.filter.packages` system property supplies values for the Ignore-Scanning-Packages property.
- See the topic on Java virtual machine custom properties.
- Add Ignore-Scanning-Archives and Ignore-Scanning-Packages entries to the application manifest, `META-INF/MANIFEST.MF`.

Note: When updating the application manifest, follow line-length limitations and other constraints for the manifest.

- Add Ignore-Scanning-Archives and Ignore-Scanning-Packages entries to the module manifest.

Note: When updating the module manifest, follow line-length limitations and other constraints for the manifest.

What to do next

Install the application again. If deployment continues to be slow, specify more modules and Java packages to ignore.

A client program does not work

What kind of problem are you seeing?

ActiveX client fails to display ASP files, or WebSphere Application Server resources (JSP files, servlet, or HTML pages) or both

A possible cause of this problem is that both IIS for serving Active Server Pages (ASP) files and an HTTP server that supports WebSphere Application Server (such as IBM HTTP Server) are deployed on the same host. This deployment leads to misdirected HTTP traffic if both servers are listening on the same port (such as the default port 80).

To resolve this problem, either:

- Open the IIS administrative panel, and edit the properties of the default web server to change the port number to a value other than 80
- Install IIS and the HTTP server on separate servers.

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

Plants by WebSphere Catalog Manager (pbwsCatalogMgr) exceptions

When you federate a stand-alone server into a Deployment Manager cell, the bootstrap port number of the application server may change. This will cause the client to not be able to communicate with the server, thus causing an exception. The following scenario may cause an exception when you start Plants by WebSphere:

1. Install a stand-alone WebSphere Application Server.
2. Run the Plants by WebSphere example.
3. Federate the stand-alone WebSphere Application Server into a Deployment Manager cell using the **addNode** command.
4. Start **pbwsCatalogMgr**.

To avoid the exception, locate the new (changed) port number on the server and modify the client configuration to match the port number on the server.

1. Go to `was_server_root\profiles\your_server_name\config\cells\your_cell\nodes\your_node`.
 - a. Open the `serverindex.xml` file.
 - b. Locate the `BOOTSTRAP_ADDRESS` port number of the application server, for example 9810.
2. Assign this port number to the client to communicate with your newly-federated application server. Go to `was_client_root\bin` and edit the `setupClient.bat` file.
3. Locate the line 'SET SERVERPORTNUMBER' and set the value for it to 9810.

If you have security enabled, ensure that the bus security is also enabled and that a user is defined to the bus connector role before running `pbwsCatalogMgr`.
4. Restart the node agent and the application server.

The client is now properly set up to start **pbwsCatalogMgr**.

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.

Web resource is not displayed

Use this information to troubleshoot problems that occur when attempting to display a resource in a browser.

If you are not able to display a resource in your browser, follow these steps:

1. Verify that your HTTP server is healthy by accessing the URL `http://server_name` from a browser and seeing whether the Welcome page appears. This action indicates whether the HTTP server is up and running, regardless of the state of WebSphere Application Server.

2. If the HTTP server Welcome page does not appear, that is, if you get a browser message like page cannot be displayed or something similar, try to diagnose your web server problem.
3. If the HTTP server appears to function correctly, the Application Server might not be serving the target resource. Try to access the resource directly through the Application Server instead of through the HTTP server.

If you cannot access the resource directly through the Application Server, verify that the URL used to access the resource is correct.

If the URL is incorrect and it is created as a link from another JavaServer Pages (JSP) file, servlet, or HTML file, try correcting it in the browser URL field and reloading, to confirm that the problem is a malformed URL. Correct the URL in the "from" HTML file, servlet or JSP file.

If the URL appears to be correct, but you cannot access the resource directly through the Application Server, verify the health of the hosting application server and web module:

- a. View the hosting application server and web module in the administrative console to verify that they are up and running.
- b. Copy a simple HTML or JSP file, such as `SimpleJsp.jsp`, which is in the WebSphere Application Server directory structure, to your web module document root, and try to access the file. If successful, the problem is with the resource.

View the logs of your Application Server to find out why your resource cannot be found or served .

4. If you can access the resource directly through the Application Server, but not through an HTTP server, the problem lies with the HTTP plug-in, the component that communicates between the HTTP server and the WebSphere Application Server.
5. If the JSP file and the servlet output are served, but not static resources such as `.html` and image files, see the steps for enabling file serving.
6. If certain resources display correctly, but you cannot display a servlet by its class name:
 - Verify that the servlet is in a directory in the web module class path, such as in the `/web_module_name.war/WEB-INF/classes` directory.
 - Verify that you specify the full class name of the servlet, including its package name, in the URL.
 - Verify that `"/servlet"` precedes the class name in the URL. For example, if the root context of a web module is `"myapp"`, and the servlet is `com.mycom.welcomeServlet`, then the URL reads:
`http://hostname/myapp/servlet/com.mycom.welcomeServlet`
 - Verify that serving the servlets by class name is enabled for the hosting web module by opening the source web module in an assembly tool and browsing the *serve servlets by classname* setting in the IBM Extensions property page. If necessary, enable this flag and redeploy the web module. For more information about the assembly tool, refer to the assembly tools section of the *Developing and deploying applications* PDF book.
 - For servlets or other resources served by mapped URLs, the URL is `http://hostname/Web module context root/mappedURL`.

Diagnosing web server problems

If you are unable to view the welcome page of your HTTP server, determine if the server is operating properly.

If the HTTP server does not start:

- Examine the HTTP server error log for clues.
- Try restoring the HTTP server to its configuration prior to installing WebSphere Application Server and restarting it. If you are using IBM HTTP Server:
 - Rename the file `IHS_install_dir\httpd.conf`.
 - Copy the `httpd.conf.default` file to the `httpd.conf` directory.
 - If Apache is running, stop and restart it.
- For the Sun ONE (iPlanet) Web Server, restore the `obj.conf` configuration file for Sun ONE V4.1 and both `obj.conf` and `magnus.conf` files for Sun ONE V6.0 and later.
- For the Microsoft Internet Information Server (IIS), remove the WebSphere Application Server plug-in through the IIS administrative GUI.

If restoring the HTTP server default configuration file works, manually review the configuration file that has WebSphere Application Server updates to verify directory and file names for WebSphere Application Server files. If you cannot manually correct the configuration, you can uninstall and reinstall WebSphere Application Server to create a clean HTTP configuration file.

If restoring the default configuration file does not help, contact technical support for the web server you are using. If you are using IBM HTTP Server with WebSphere Application Server, check available online support (hints and tips, technotes, and fixes). If you do not find your problem listed there, see Troubleshooting help from IBM.

Accessing a web resource through the application server and bypassing the HTTP server

You can bypass the HTTP server and access a web resource through the application server. It is not recommended to serve a production website in this way, but it provides a good diagnostic tool when it is not clear whether a problem resides in the HTTP server, WebSphere Application Server, or the HTTP plug-in.

To access a web resource through the Application Server:

1. Determine the port of the HTTP service in the target application server.
 - a. In the WebSphere administrative console, click **Servers > Application Servers**.
 - b. Select the target server, then under Additional Properties click **Web Container**.
 - c. Under the Additional Properties of the web container, click **HTTP Transports**. You see the ports listed for virtual hosts served by the application server.
2. Use the HTTP transport port number of the application server to access the resource from a browser. For example, if the port is 9080, the URL is `http://hostname:9080/myAppContext/myJSP.jsp`.
3. If you are still unable to access the resource, verify that the HTTP transport port is in the "Host Alias" list:
 - a. Click **Application Servers > Your_ApplicationServer > Web Container > HTTP Transports** to check the Default virtual host and the HTTP transport ports used by this application server.
 - b. Click **Environment > Virtual Hosts > default host > Host Aliases** to check if the HTTP transport port exists. Add an entry if necessary. For example, if the HTTP port for your application is server is 9080, add a host alias of `*:9082`.

Application uninstallation problems

When you try to uninstall an application or node, you might encounter problems. This topic suggests ways to resolve uninstallation problems.

What kind of problem are you having?

- After uninstalling an application through wsadmin tool, the application continues to run and throws "DocumentIOException"
- The removeNode command does not remove the installed application from the deployment manager
- I cannot display the syntax for the removeNode command.

If none of these steps fixes your problem:

- Make sure that the application and its web and EJB modules are in a stopped state before uninstalling.
- If you are uninstalling or installing an application using wsadmin, make sure that you are using the `-conntype NONE` option to invoke wsadmin and enable local mode. To use the `-conntype NONE` option, stop the hosting application server before uninstalling the application.
- Check to see if the problem has been identified and documented by looking at the available online support (hints and tips, technotes, and fixes).
- If you don't find your problem listed there, contact IBM support

After uninstalling application through the wsadmin tool, the application throws "DocumentIOException"

If this exception occurs after the application was uninstalled using wsadmin with the `-conntype NONE` option:

- Restart the server or,
- Rerun the uninstall command without the `-conntype NONE` option.

The removeNode command does not remove the installed application from the deployment manager

If the applications were installed indirectly using the `addNode` command with the `-includeapps` option, then `removeNode` will not uninstall them, since they may be in use by other nodes. These applications must be explicitly uninstalled, for example through the administrative console.

I cannot display the syntax for the removeNode command

Unlike the `addNode` command, the `removeNode` command is valid with no parameters, so executing it will execute the operation, that is, remove the node, without displaying the command syntax.

To see the valid options for `removeNode`, run `removeNode -?` or `removeNode -help`.

Chapter 14. Troubleshooting administration

Use this information if you are having problems with administrative functions.

Procedure

- Select the problem you are experiencing.
 - I have problems bringing up or using the administrative console.
 - I have problems starting or using the **wsadmin** command prompt.
 - My web module or application server dies or hangs.
 - I get errors trying to configure and enable security.
 - I have problems creating or using HTTP sessions.
 - I have problems using tracing, logging, log files, or other troubleshooting features.
 - I get errors connecting to the administrative console from a browser.
- If you did not solve the problem, prepare to contact IBM support.

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

If you do not see a problem that resembles yours, or if the information provided does not solve your problem, contact IBM support for further assistance.

Administration and administrative console troubleshooting

Note: This topic references one or more of the application server log files. Beginning in WebSphere Application Server Version 8.0 you can configure the server to use the High Performance Extensible Logging (HPEL) log and trace infrastructure instead of using SystemOut.log, SystemErr.log, trace.log, and activity.log files or native z/OS logging facilities. If you are using HPEL, you can access all of your log and trace information using the LogViewer command-line tool from your server profile bin directory. See the information about using HPEL to troubleshoot applications for more information on using HPEL.

In WebSphere Application Server products, administrative functions are supported by:

- The application server (such as server1) process
- The deployment manager process in the WebSphere Application Server, Network Deployment product

The process must be running to use the administrative console. The **wsadmin** command-line utility has a local mode that you can use to perform some administrative functions, even when the server process is not running.

What kind of problem are you seeing?

- “Server status and messages in the console view are not current”
- “Role-based authorization fails” on page 578
- “When starting or stopping a server using a wsadmin interactive scripting session, you receive an exception indicating read timed out” on page 579
- “Problems starting or using the administrative console or wsadmin utility” on page 579

Server status and messages in the console view are not current

When connecting to an Application Server that uses a Simple Object Access Protocol (SOAP) connection for a long time, the following problems begin to occur:

- Under the status column in the Servers view on an administrative console panel, the status of the server does not refresh.
- Server messages are not updated in the administrative console.
- A decrease of system resources occurs as numerous ports are created and left in the TIME_WAIT state.

This problem persists even after you restart the server or you start another server that uses the SOAP connection port.

The problem occurs because the SOAP connector does not support connection pooling. If the Application Server has many ongoing operations that use the SOAP connector, the Application Server quickly opens and closes many ports. Due to the nature of the underlying TCP/IP protocol, these ports remain in the TIME_WAIT state for some time before the operating system can reclaim them. The number of ports that WebSphere Application Server opens can exceed the limit that the operating system imposes. Under this condition, the opening of additional ports fails through the SOAP connector until the operating system reclaims ports.

Use the following options to work around the problem:

- Increase the operating system limits on the number of ports.
- For Rational Application Developer, the wsadmin utility, or Java applications that use the Java Management Extension (JMX) connectors, switch to the Remote Method Invocation (RMI) connector.
- Wait until few or no ports are in the TIME_WAIT state before performing new operations through Rational Application Developer or the administrative console.

Role-based authorization fails

When you make a Java Management Extension (JMX) call such as `getAttribute`, `setAttribute`, `invoke`, and so on in your application, the caller requires an administrative role with sufficient permissions. The required role depends on the MBean attribute or method that the JMX caller calls and can be one of administrator, configurator, monitor, or operator. If one of the administrative roles is not assigned to the caller, or if the role is assigned, but the caller does not have the required permissions, the application receives a role-based authorization failure, for example:

```
SECJ0305I: Role based authorization check failed for securityname server.domain.name:3890/user.id,
accessId user:server.domain.name:3890/uid=user.id,ou=xxxx,dc=yyy,dc=zzz while invoking method
getNodeName on resource Server and module Server.
```

If the caller of the application cannot be assigned one of the administrative roles, the application can log in with one of the roles on behalf of the caller. For example:

```
try
{
    // Create a LoginContext to authenticate a user ID and password.
    javax.security.auth.login.LoginContext
    lc = new javax.security.auth.login.LoginContext("WSLogin",
    new com.ibm.websphere.security.auth.callback.WSCallbackHandlerImpl("adminuser",
    "adminpassword"));

    // perform the login
    lc.login();

    // Get the authenticated subject.
    javax.security.auth.Subject adminSubject = lc.getSubject();

    // Define the action that will take place using the authenticated Subject
    // You can define this action anywhere in the code, the action
    // is reference in the WSSubject.doAs that follows.
    java.security.PrivilegedAction adminAction = new java.security.PrivilegedAction()
    {
        public Object run()
        {
            try
            {
                // Get the WebSphere AdminService.
                AdminService adminservice = AdminServiceFactory.getAdminService();

                // Get the WebSphere Admin Local Server MBean instance.
                ObjectName objectname = adminservice.getLocalServer();
```

```

// Get the Node name.
String nodeName = (String)adminservice.getAttribute(objectname, "nodeName");

// Get the Application Server name.
String serverName = (String)adminservice.getAttribute(objectname, "name");

// Get the Application Server Process ID.
String serverPid = (String)adminservice.getAttribute(objectname, "pid");

// Return a result, for this example, just return the serverPid.
return serverPid;
}
catch (Exception e)
{
    e.printStackTrace();
}
return null;
}
});

// Invoke an AdminClient resource using the authenticated subject.
// This example demonstrates the action of creating an
// administrative client and returning a String value to use outside
// the doAs block.
String myData = (String)
    com.ibm.websphere.security.auth.WSSubject.doAs(adminSubject, adminAction);

// use "myData" later on....
}
catch (javax.security.auth.login.LoginException e)
{
    e.printStackTrace();
}
}

```

When starting or stopping a server using a wsadmin interactive scripting session, you receive an exception indicating read timed out

When starting or stopping a server using a wsadmin interactive scripting session, you receive an exception indicating read timed out, for example:

```

WASX7015E: Exception running command: "$AdminControl startServer server1 Node1";
exception information: com.ibm.websphere.management.exception.ConnectorException
org.apache.soap.SOAPException: [SOAPException: faultCode=SOAP-ENV:Client; msg=Read
timed out; targetException=java.net.SocketTimeoutException: Read timed out]

```

This exception occurs because the timeout value is too small. Increase the timeout value specified by the `com.ibm.SOAP.requestTimeout` property in the `soap.client.props` file in the `profile_root/properties` directory for a single server edition or in the `profile_root/properties` directory for a network deployment installation. The value you choose depends on a number of factors such as the size and the number of the applications installed on the server, the speed of your machine, and the level of usage of your machine. The default value of the `com.ibm.SOAP.requestTimeout` property is 180 seconds.

Problems starting or using the administrative console or wsadmin utility

If you have problems starting or using the administrative console or wsadmin utility, verify that the supporting server process is started and that it is healthy.

- For the application server process, look at these files:
 - `profile_root/logs/server_name/startServer.log` for the message that indicates that the server started successfully: `ADMU3000I: Server server1 open for e-business; process id is nnnn..`
 - The server log files for the message that indicates that the server started successfully: `WSVR0001I: Server server open for e-business.`
- For the WebSphere Application Server, Network Deployment product, look at these files:

- *profile_root*/logs/dmgr/startServer.log for the message that indicates that the server started successfully: ADMU3000I: Server dmgr open for e-business; process id is *nnnn*.
- The server log files for this message indicating that the server started successfully: WSVR0001I: Server dmgr open for e-business.
- For the z/OS product, check the job output.
- Look up any error messages in these files in the message reference table. Select the **Reference** view in the information center navigation, and click **Messages**. A message like WASX7213I: This scripting client is not connected to a server process when trying to start wsadmin indicates that either the server process is not running, the host machine where it is running is not accessible, or that the port or server name that the wsadmin utility uses is incorrect.
- Verify that you are using the right port number to communicate with the administrative console or the wsadmin server:
 - Look in the joblogs file.
 - The line ADMC0013I: SOAP connector available at port *nnnn* indicates the port that the server is using to listen for wsadmin functions.
 - The com.ibm.ws.scripting.port property in the *profile_root*/properties/wsadmin.properties file controls the port used by the wsadmin utility to send requests to the server.
 - If port value is different from the value shown in the SystemOut.log file, either change the port number in the wsadmin.properties file, or specify the correct port number when starting the wsadmin utility by using the -port *port_number* property on the command line.

The com.ibm.ws.scripting.port property in the *profile_root*/properties/wsadmin.properties file controls the port used by the wsadmin utility to send requests to the server.
 - If the port value is different than the one specified in the web address for the administrative console, change the web address in the browser to the correct value. The default value is <http://localhost:9060/ibm/console>.
- Use the **telnet** command to test that the host name where the application server or deployment manager is running, is reachable from the system where the browser or wsadmin program is used. If you can ping the host name, no firewall or connectivity issues exist.
- If the host where the application server or deployment manager is running is remote to the machine from which the client browser or wsadmin command is running, ensure that the appropriate host name parameter is correct. Verify:
 - The host name in the browser web address for the console.
 - The -host *host name* option of the **wsadmin** command that is used to direct the wsadmin utility to the right server

If none of these steps solves the problem, see if the specific problem you are having is addressed in the Installation completes but the administrative console does not start topic. Check to see if the problem has been identified and documented using the links in the Diagnosing and fixing problems: Resources for learning topic. If you do not see a problem that resembles yours, or if the information provided does not solve your problem, contact IBM support for further assistance.

For current information available from IBM Support on known problems and their resolution, see the following topics on the IBM support page:

- Administrative Console
- Administrative Scripting Tools
- System management

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

- Administrative Console
- Administrative Scripting Tools
- System management

Administrative console does not start even though installation completes

This topic discusses problems that you can encounter when you attempt to access the console.

Note: This topic references one or more of the application server log files. Beginning in WebSphere Application Server Version 8.0 you can configure the server to use the High Performance Extensible Logging (HPEL) log and trace infrastructure instead of using `SystemOut.log`, `SystemErr.log`, `trace.log`, and `activity.log` files or native z/OS logging facilities. If you are using HPEL, you can access all of your log and trace information using the LogViewer command-line tool from your server profile bin directory. See the information about using HPEL to troubleshoot applications for more information on using HPEL.

What kind of problem are you having?

- “An Internal Server Error, Page cannot be found, 404, or similar error occurs trying to view the administrative console”
- “An Unable to process login. Check user ID and password and try again. error occurs when trying to access the administrative console page” on page 582
- “The directory paths in the administrative console contain strange characters” on page 582

If you can bring up the browser page, but the administrative console behavior is inconsistent, error prone, or unresponsive, try upgrading your browser. Older browsers might not support all the features of the administrative console. For a listing of supported web browsers, see the Supported hardware and software web page.

IBM Support has documents and tools that can save you time gathering information needed to resolve problems as described in Troubleshooting help from IBM. Before opening a problem report, see the Support page:

- http://www.ibm.com/software/webservers/appserv/zos_os390/support/

An Internal Server Error, Page cannot be found, 404, or similar error occurs trying to view the administrative console

Here are some steps to try if you are unable to view the administrative console:

- Verify that the application server which supports the administrative console is up and running.
 - For a base configuration, the administrative console is deployed by default on server1. Before viewing the administrative console, you must verify that server1 is running. Do so by issuing the following command on the MVS console to list active processes:

```
D A,L
```

Note: See z/OS MVS System Commands for information on how to use MVS operator commands. Check for the application server procedure name of BB05ACR with the server short name of BB0S001. If it is not running, enter the following command on the MVS console:

```
START appserver_proc_name,JOBNAME=server_short_name,  
      ENV=cell_short_name.node_short_name.server_short_name
```

Example:

```
START BB06ACR,JOBNAME=BB0S001,ENV=PLEX1.SY1.BB0S001
```

- For a network deployment configuration, the deployment manager runs the administrative console. For example, to start the deployment manager, you can issue the following command from the MVS console:

```
START BB05DCR,JOBNAME=BB0DMGR,ENV=PLEX1.PLEX1.BB0DMGR
```
 - View the joblog or stdout for the application server or deployment manager to verify that the server supporting the administrative console has started.
- Check the web address you use to view the console.

- Try to eliminate connection, address and firewall issues by pinging the server machine from a command prompt, using the server name in the URL.

An Unable to process login. Check user ID and password and try again. error occurs when trying to access the administrative console page

This error indicates that security is enabled for WebSphere Application Server, and that the user ID or password supplied is either not valid or not authorized to access the console.

To access the console:

- If you are the administrator, use the ID defined as the security administrative ID. This ID is stored in the WebSphere Application Server `security.xml` file.

The directory paths in the administrative console contain strange characters

Directory paths that are used for class paths or resources specified in an assembly tool, in configuration files, or elsewhere that contain strange characters when they are viewed in the administrative console might result from the Java run time interpreting a backslash (\) as a control character.

To resolve, make sure you have no back slashes in your class paths. Replace any back slashes with forward slashes.

Administrative console - browser connection problems

This topic describes problems that you can have when logging into the administrative console from a browser.

Review the following information to resolve your browser problem.

If you are able to bring up the browser page, but the console behavior is inconsistent, error-prone, or unresponsive, try upgrading the browser you are using. Older browsers may not support the administrative console's features. For a listing of supported web browsers, see the Supported hardware and software web page.

Check to see if the problem has been identified and documented using the links in Diagnosing and fixing problems: Resources for learning. If you do not see a problem that resembles yours, or if the information provided does not solve your problem, contact IBM support for further assistance.

Check the following list for your problem and how to solve it:

- When a single user that uses multiple instances of the Mozilla browser logs into the administrative console, the first user ID that logs into the administrative console is the current user.
- A user on Mozilla browser Version 1.4 selects a check box on a collection table, presses Enter, and receives an error.
- A user on Mozilla browser Version 1.4 enters an invalid ID or password, presses Enter, and receives an error message

When a single user that uses multiple instances of the Mozilla browser logs into the administrative console, the first user ID that logs into the administrative console is the current user.

When a single user logged into an operating system tries to use multiple instances of the Mozilla browser, the first user ID that logs into the administrative console is the current user. This situation occurs because the browser windows share a single process.

To resolve the problem, do one of the following actions:

- Have single users logged into an operating system use a single instance of the Mozilla browser to log into the administrative console.
- If the operating system allows multiple users on an operating system, have each user log into the operating system with a different user ID and bring up a single instance of the Mozilla browser.

A user on Mozilla browser Version 1.4 selects a check box on a collection table, presses Enter, and receives an error.

A user on Mozilla browser Version 1.4 selects a check box on a collection table, presses Enter, and receives an error.

To resolve the problem, do one of the following actions:

- Explicitly select a button of interest on the administrative console panel instead of pressing Enter.
- Use a later version of a supported Mozilla browser.
- Use a supported version of the Microsoft Internet Explorer browser.

A user on Mozilla browser Version 1.4 enters an invalid ID or password, presses Enter, and receives an error message

A user on Mozilla browser Version 1.4 enters an invalid user ID or password, presses Enter, and receives an error message. Clicking OK fails to refresh the administrative login screen

To resolve the problem, do one of the following actions:

- Use a later version of a supported Mozilla browser.
- Use a supported version of the Microsoft Internet Explorer browser.

Web server plug-in troubleshooting tips

The following topics might help you diagnose problems with the web server plug-ins.

If you are having problems using a web server plug-in, try these steps:

- Review the file *plugins_root/logs/web_server_name/http_plugin.log* for clues. Look up any error or warning messages in the message table.
- Review your web server's error and access logs to see if the web server is having a problem:
 - IBM HTTP Server and Apache: *access.log* and *error.log*.
 - Domino Web Server: *httpd-log* and *httpd-error*.
 - Sun Java System: *access* and *error*.
 - Microsoft IIS: *timedatestamp.log*.

If these files don't reveal the cause of the problem, follow these additional steps.

Plug-in Problem Determination Steps

The plug-in provides very readable tracing which can be beneficial in helping to figure out the problem. By setting the LogLevel attribute in the *config/plugin-cfg.xml* file to **Trace**, you can follow the request processing to see what is going wrong.

Note: If you are using a Veritas File System with large file support enabled, file sizes up to two terabytes are allowed. In this case, if you set the LogLevel attribute in the *plugin-cfg.xml* file to LogLevel=Trace, then the *http_plugin.log* file might grow quickly and consume all available space on your file system. Therefore, you should set the value of the LogLevel attribute to ERROR or DEBUG to prevent high CPU utilization..

At a high level, complete these steps.

1. The plug-in gets a request.

2. The plug-in checks the routes defined in the `plugin-cfg.xml` file.
3. It finds the server group.
4. It finds the server.
5. It picks the transport protocol, HTTP or HTTPS.
6. It sends the request.
7. It reads the response.
8. It writes it back to the client.

You can see this very clearly by reading through the trace for a single request:

- The first step is to determine if the plug-in has successfully loaded into the web server.
 - Check to make sure the `http_plugin.log` file has been created.
 - If it has, look in it to see if any error messages indicate some sort of failure that took place during plug-in initialization. If no errors are found look for the following stanza, which indicates that the plug-in started normally. Ensure that the timestamps for the messages correspond to the time you started the web server:

```
[Thu Jul 11 10:59:15 2002] 000009e 000000b1 - PLUGIN: -----System Information-----
[Thu Jul 11 10:59:15 2002] 000009e 000000b1 - PLUGIN: Bld date: Jul 3 2002, 15:35:09
[Thu Jul 11 10:59:15 2002] 000009e 000000b1 - PLUGIN: Web server: IIS
[Thu Jul 11 10:59:15 2002] 000009e 000000b1 - PLUGIN: Hostname = SWEETTJ05
[Thu Jul 11 10:59:15 2002] 000009e 000000b1 - PLUGIN: OS version 4.0, build 1381, 'Service Pack 6'
[Thu Jul 11 10:59:15 2002] 000009e 000000b1 - PLUGIN: -----
```

- Some common errors are:

lib_security: loadSecurityLibrary: Failed to load gsk library

Either GSKit did not get installed or the wrong version of GSKit got installed. To determine which situation occurred:

If you cannot find the appropriate file, try reinstalling the plug-in with the correct GSKit version to see if this fixes the problem.

ws_transport: transportInitializeSecurity: Keyring wasn't set

The HTTPS transport defined in the configuration file was prematurely terminated and did not contain the Property definitions for the keyring and stashfile. Check your XML syntax for the line number given in the error messages that follow this one to make sure the Transport element contains definitions for the keyring and stashfiles before it is terminated.

- If the `http_plugin.log` file is not created, check the web server error log to see if any plug-in related error messages have been logged there that indicate why the plug-in is failing to load. Typical causes of this can include failing to correctly configure the plug-in with the web server environment. Check the documentation for configuring the web server that you are using with the web server plug-in.

- Determine whether there are network connection problems with the plug-in and the various application servers defined in the configuration. Typically you will see the following message when this is the case:

ws_common: webspHEREGetStream: Failed to connect to app server, OS err=%d

Where %d is an OS specific error code related to why the `connect()` call failed. This can happen for a variety of reasons.

- Ping the machines to make sure they are properly connected to the network. If the machines cannot be pinged, there is no way for the plug-in to contact them. Possible reasons for this problem include:
 - Firewall policies that limit the traffic from the plug-in to the application server.
 - The machines are not on the same network.
- If you are able to ping the machines then the probable cause of the problem is that the port is not active. The port might not be active because the application server or cluster is not started or the application server has gone down for some reason. To verify that this is the problem, you can try to manually telnet into the port that the connect is failing on. If you cannot telnet into the port the application server is not up and that problem needs to be resolved before the plug-in can successfully connect.
- Determine whether other activity on the machines where the servers are installed is impairing the ability of the server to service a request. Check the processor utilization as measured by the task manager, processor ID, or some other outside tool to see if it:
 - Is not what was expected.
 - Is erratic rather than a constant.

- Shows that a newly added member of the cluster is not being utilized.
- Shows that a failing member that has been fixed is not being utilized.
- Check the administrative console for server status.
Check the administrative console to ensure that the application servers are started. View the administrative console for error messages.
- In the administrative console, select the problem application server and view its installed applications to verify that they are started.

If none of these steps solves the problem:

- For specific problems that can cause web pages and their contents not to display, see Web resource (JSP file, servlet, html file, image, etc) will not display in the information center.
- Check to see if the problem has been identified and documented using the links in Diagnosing and fixing problems: Resources for learning.
- If you do not see a problem that resembles yours, or if the information provided does not solve your problem, contact IBM support for further assistance.

For current information available from IBM Support on known problems and their resolution, see the following topics on the IBM support page:

- HTTP transport
- HTTP plug-in
- HTTP plug-in remote install

For current information available from IBM Support on known problems and their resolution, see the IBM Support page. You should also refer to this page before opening a PMR because it contains documents that can save you time gathering information needed to resolve a problem.

You might find the following topics on the IBM support page helpful:

- HTTP plug-in
- HTTP plug-in remote install

Multiserver environment errors

Use this information to troubleshoot problems with setting up multiserver environments.

What kind of problem are you seeing?

- “After creating and starting a cluster, the cluster does not start, and logs show that servers in the cluster are not found” on page 586
- “One or more nodes do not display in the administrative console” on page 586
- “The addNode command fails” on page 587
- “Application files are not present on all nodes” on page 587
- “After downloading the WebSphere Application Server, Network Deployment plug-in to my system, my server does not start” on page 587
- “In a clustered environment, a server with debug mode enabled does not start” on page 587

If none of these problem solution descriptions fixes your problem:

1. Browse the logs of the problem deployment manager and applications servers.
 - a. Look up any error messages by selecting the **Reference** view of the information center navigation and expanding **Messages** in the navigation tree.
 - b. If Java exceptions are displayed in the log files, try to determine the actual subcomponent directly involved in the problem by examining the trace stack and looking for a WebSphere Application Server-related class near the top of the stack (names beginning with `com.ibm.websphere` or `com.ibm.ws`) that created the exception.

For example, if the exception seems to be created by a class in the `com.ibm.websphere.naming` package, review the Naming services component troubleshooting tips topic.

2. Ensure that all the machines in your configuration have TCP/IP connectivity to each other by running the **ping** command:
 - a. From each physical server to the deployment manager
 - b. From the deployment manager to each physical server
3. Although the problem is occurring in a clustered environment, the actual cause might be indirectly related, or unrelated to clustering. Investigate all relevant possibilities:
 - a. If an enterprise bean on one or more servers is not serving requests, review the Enterprise bean cannot be accessed from a servlet, a JSP file, a stand-alone program, or another client and Application access problems topics.
 - b. If problems seem to occur after enabling security, review the Access problems after enabling security topic.
 - c. If an application server stops responding to request, or spontaneously fails (its process closes), review the Web module or application server stops processing requests topic.
 - d. If SOAP requests are not served by some servers, review the Application client SOAP request troubleshooting tips topic.
 - e. If you have problems installing or deploying an application on servers on one or more nodes, review the “Application deployment problems” on page 557 topic.
4. If your topology consists of a Windows-based deployment manager with UNIX-based servers, browse any recently-updated `.xml` and `.policy` files on the UNIX-based platform using the vi editor to ensure that Control-M characters are not present in the files. Edit these files using the vi editor on the UNIX-based platform, to avoid inserting these characters.
5. Check the steps for troubleshooting the workload management component.
6. Check to see if the problem is identified and documented by looking at available online support (hints and tips, technotes, and fixes).

After creating and starting a cluster, the cluster does not start, and logs show that servers in the cluster are not found

This error can occur when the configuration is not synchronized from the deployment manager to a node. If **auto synchronization** is enabled, wait until the synchronization has run. If you are using manual synchronization, explicitly request a synchronization to each node on the cluster.

To determine whether synchronization has occurred, look at the configuration on the node machines using the administrative console and verify that the new cluster members are defined on each node.

One or more nodes do not display in the administrative console

This problem can occur when a basic connectivity problem exists between the deployment manager server and other servers in the topology. Look for the `serverindex.xml` file in the deployment manager directory structure:

- If the problem node does not display in the list, review the steps for adding a node to the cluster.
- If the problem node does display in the list:
 - From the deployment manager server, ping the server name as it displays in the list. If the **ping** command indicates no communication, verify that the host name is correct in the list, correct it if necessary, then restart the deployment manager.
 - If the name that displays in the list is the short name, ping the fully qualified network name. If the corrected name works, update the list, and restart the deployment manager.
 - If the problem server uses Dynamic Host Configuration Protocol (DHCP), try replacing the logical host name with the IP address and restart the deployment manager. If this action resolves the problem, be aware that you must change the `serverindex.xml` file each time the problem server address changes, and potentially each time the problem machine is rebooted. To avoid this problem, consider assigning a static IP address to the server.
- If you still cannot establish communication between the servers, contact your network administrator to resolve the problem, and restart the deployment manager after the problem is corrected.

The addNode command fails

This error can occur when the deployment manager Domain Name Server (DNS) configuration is set up improperly. The default installation on Linux systems uses the loopback address (127.0.0.1) as the default host address. To verify this problem, query the host name of the suspect machine. If the query returns localhost 127.0.0.1, or if the file transfer traces at the node show that the node is trying to upload files to a web address that includes 127.0.0.1, the node has an incorrect DNS configuration.

To correct this problem, update the `/etc/hosts` file or the name service configuration file, `/etc/nsswitch.conf`, to query the Domain Name Server or Network Information Server (NIS) before searching hosts.

Application files are not present on all nodes

In the WebSphere Application Server, Network Deployment environment, application binary files are transferred to the individual nodes where applications are supported as part of the node synchronization operation. During node synchronization, application files are only propagated if their deployment descriptors specify `enableDistribution=true`. This flag is specified as part of the application installation procedure in the administrative console, and is stored as a property in the `app_server_root/config/cells/cell_name/applications/application_name/deployment.xml` file.

To confirm this problem, check to see whether the `enableDistribution` flag is set. If it is already set to true, ensure that the target node is configured to run auto file synchronization.

If both of these settings are correct and the problem persists, manually perform a synchronization. If the application files still do not display in the installation directory, use the EARExpander tool in `app_server_root/bin` directory to expand the EAR file from the repository to the installation destination. On remote nodes, the repository displays in the `config/cells/cell_name/applications/application_name.ear/` directory.

After downloading the WebSphere Application Server, Network Deployment plug-in to my system, my server does not start

If you experience this situation, the most likely cause is that the transport paths in the plug-in must be modified to work in your environment. See the Example: Manually editing transport settings in the `server.xml` file topic for information on how to modify these settings.

In a clustered environment, a server with debug mode enabled does not start

This problem occurs when the following three conditions exist:

- Multiple server processes are configured to run on the same node
- More than one server has debug mode enabled
- The debug arguments for more than one of the servers are left at the default values, so that more than one server in the node is trying to use the same debug port (port number 7777).

The server does not start because multiple servers processes running on the same physical host machine with debug enabled cannot use the same debug port.

To correct this problem, for each server:

1. On the administrative console click **Server > Application servers > *server_name* > Java and Process Management > Process Definition > Java Virtual Machine**.
2. Update the debug argument so that the address of the debug port (`address=port number`) is unique for each server process.

Workload management component troubleshooting tips

If the workload management component is not properly distributing the workload across servers in multi-node configuration, use the following options to isolate the problem.

Note: This topic references one or more of the application server log files. Beginning in WebSphere Application Server Version 8.0 you can configure the server to use the High Performance Extensible Logging (HPEL) log and trace infrastructure instead of using `SystemOut.log`, `SystemErr.log`, `trace.log`, and `activity.log` files or native z/OS logging facilities. If you are using HPEL, you can access all of your log and trace information using the LogViewer command-line tool from your server profile bin directory. See the information about using HPEL to troubleshoot applications for more information on using HPEL.

- Ensure that the workload is distributed across clustered servers.
- Resolve any problems with the multiserver Deployment Manager environment setup.
- “Eliminate environment or configuration issues”
- “Resolve problem or contact IBM support” on page 589

Eliminate environment or configuration issues

Determine if the servers are capable of serving the applications for which they have been enabled. Identify the cluster that has the problem.

- Are there network connection problems with the members of the cluster or the administrative servers, for example deployment manager or node agents?
 - If so, ping the machines to ensure that they are properly connected to the network.
- Is there other activity on the machines where the servers are installed that is impacting the servers ability to service a request? For example, check the processor utilization as measured by the task manager, processor ID, or some other outside tool to see if:
 - It is not what is expected, or is erratic rather than constant.
 - It shows that a newly added, installed, or upgraded member of the cluster is not being utilized.
- Are all of the application servers you started on each node running, or are some stopped?
- Are the applications installed and operating?
- If the problem relates to distributing workload across container-managed persistence (CMP) or bean-managed persistence (BMP) enterprise beans, have you configured the supporting JDBC providers and JDBC data source on each server?

If you are experiencing workload management problems related to HTTP requests, such as HTTP requests not being served by all members of the cluster, be aware that the HTTP plug-in balances the load across all servers that are defined in the PrimaryServers list if affinity has not been established. If you do not have a PrimaryServers list defined then the plug-in load balances across all servers that are defined in the cluster if affinity has not been established. If affinity has been established, the plug-in should go directly to that server for all requests.

For workload management problems relating to enterprise bean requests, such as enterprise bean requests not getting served by all members of a cluster:

- Are the weights set to the allowed values?
 - For the cluster in question, log onto the administrative console and:
 1. Select **Servers > Clusters > WebSphere application server clusters**.
 2. Select your cluster from the list.
 3. Select **Cluster members**.
 4. For each server in the cluster, click on *server_name* and note the assigned weight of the server.
 - Ensure that the weights are within the valid range of 0-20. If a server has a weight of 0, no requests are routed to it. Weights greater than 20 are treated as 0.

The remainder of this article deals with enterprise bean workload balancing only. For more help on diagnosing problems in distributing web (HTTP) requests, view the "Web server plug-in troubleshooting tips" and "Web resource does not display" topics.

Resolve problem or contact IBM support

If the client logs indicate an error in WLM, collect the following information and contact IBM support.

- A detailed description of your environment.
- A description of the symptoms.
- The server log files for all servers in the cluster.
- A description of what the client is attempting to do, and a description of the client. For example, 1 thread, multiple threads, servlet, J2EE client, etc.

If none of these steps solves the problem, check to see if the problem has been identified and documented using the links in the "Diagnosing and fixing problems: Resources for learning" topic. If you do not see a problem that resembles yours, or if the information provided does not solve your problem, contact IBM support for further assistance.

If you do not find your problem listed there, contact IBM Support.

Workload is not getting distributed

This information might help you diagnose the trouble if you are having a workload distribution problem.

Note: This topic references one or more of the application server log files. Beginning in WebSphere Application Server Version 8.0 you can configure the server to use the High Performance Extensible Logging (HPEL) log and trace infrastructure instead of using `SystemOut.log`, `SystemErr.log`, `trace.log`, and `activity.log` files or native z/OS logging facilities. If you are using HPEL, you can access all of your log and trace information using the LogViewer command-line tool from your server profile bin directory. See the information about using HPEL to troubleshoot applications for more information on using HPEL.

What kind of problem are you seeing?

- "HTTP requests are not distributed to all servers" on page 590
- "Enterprise bean requests are not distributed to all servers" on page 590
- "A failing server still receives enterprise bean requests (failover is not completed)" on page 590
- "Stopped or hung servers do not share the workload after being restored" on page 590

If none of these problem solution descriptions fix your problem:

1. Browse the JVM logs of the problem deployment manager and application servers:
 - a. Look up any error messages by selecting the **Reference** view of the information center navigation and expanding **Messages** in the navigation tree.
 - b. If Java exceptions appear in the log files, try to determine the actual subcomponent that is directly involved in the problem by examining the trace stack and looking for a product-related class near the top of the stack (names beginning with `com.ibm.websphere` or `com.ibm.ws`) that created the exception. If appropriate, review the steps for troubleshooting the appropriate subcomponent under the *Troubleshooting WebSphere applications* section of the Information Center.

For example, if the exception appears to have been thrown by a class in the `com.ibm.websphere.naming` package, review the Naming Services Component troubleshooting tips topic.
2. Ensure that all the machines in your configuration have TCP/IP connectivity to each other by running the **ping** command:
 - a. From each physical server to the deployment manager
 - b. From the deployment manager to each physical server
3. Although the problem is happening in a clustered environment, the actual cause might be only indirectly related, or unrelated, to clustering. Investigate all relevant possibilities:
 - a. If an enterprise bean on one or more servers is not serving requests, review the "Cannot access an enterprise bean from a servlet, JSP, stand-alone program, or other client" and "Cannot look up an object hosted by the product from a servlet, JSP file, or other client" topics.

- b. If problems seem to appear after enabling security, review the "Errors or access problems after enabling security" topic.
 - c. If an application server stops responding to requests, or spontaneously dies (its process closes), review the "Web module or application server dies or hangs" topic.
 - d. If SOAP requests are not being served by some or all servers, review the "Errors returned to client trying to send a SOAP request" topic.
4. Check to see if the problem is identified and documented by looking at available online support (hints and tips, technotes, and fixes).

HTTP requests are not distributed to all servers

If HTTP requests are not being distributed to all servers:

- Check your Primary Servers list. The plug-in load balances across all servers that are defined in the Primary Servers list, if affinity has not been established. If you do not have a Primary Servers list defined, the plug-in load balances across all servers defined in the cluster, if affinity has not been established. In the case where affinity has been established, the plug-in should go directly to that server, for all requests within the same HTTP session.
- If some servers are servicing requests and one or more others are not, try accessing a problem server directly to verify that it works, apart from workload management issues. If that does not work:
 - Use the administrative console to ensure that the affected server is running.
 - See the topic "Web resource does not display" for more information.
- See the "HTTP plug-in component troubleshooting tips" topic for more information.

Enterprise bean requests are not distributed to all servers

If a client cannot reach a server in a cluster thought to be reachable, a server might be marked unusable, or is down. To verify this:

- Use the administrative console to verify that the server is started. Try starting it, or if started, stop and restart it.
- Browse the administrative console and verify that the node that runs the server having the problem appears. If it does not:
 - Review the steps for adding a node to a cluster.
 - Review the steps in the section *One or more nodes do not show up in the administrative console*.
- If possible, try accessing the enterprise bean directly on the problem server to see if there is a problem with TCP/IP connectivity, application server health, or other problem not related to workload management. If this fails, review the "Cannot access enterprise bean from a servlet, JSP, stand-alone program , or other client" topic.

A failing server still receives enterprise bean requests (failover is not completed)

Some possible causes of this problem are:

- The client might have been in a transaction with an enterprise bean on the server that went down. Check the JVM logs of the application server hosting the problem enterprise bean instance. If a request is returned with **CORBA SystemException COMM_FAILURE org.omg.CORBA.completion_status.COMPLETED_MAYBE**, this might be working as designed. The design is to let this particular exception flow back to the client, since the transaction might have completed. Failing over this request to another server could result in this request being serviced twice.
- If the requests sent to the servers come back to the client with any other exceptions consistently, it might be that no servers are available.

Stopped or hung servers do not share the workload after being restored

This error occurs when previously unavailable servers are not recognized by the workload management component after those servers are restored. There is an **unusable** interval determined by the property `com.ibm.websphere.wlm.unusable.interval` during which the workload manager waits to send to a server that has been marked unusable. By default this is 5 minutes.

You can confirm that this is the problem by ensuring that servers that were down are now up and capable of servicing requests. Then wait for the unusable interval to elapse before checking to determine whether failover occurs.

Administrative problems with the wsadmin scripting tool

Use this information if you are having problems starting or using the wsadmin tool.

What kind of problem are you having?

- WASX7016E, WASX7017E, or WASX7209I: Jython scripting language error
- "WASX7023E: Error creating "SOAP" connection to host" or similar error trying to launch wsadmin command line utility.
- "com.ibm.bsf.BSFException: error while evaluating Jacl expression: no such method "<command name>" in class com.ibm.ws.scripting.AdminConfigClient" returned from wsadmin command.
- WASX7022E returned from running "wsadmin -c ..." command, indicating invalid command.
- com.ibm.ws.scripting.ScriptingException: WASX7025E: String "" is malformed; cannot create ObjectName.
- WASX701E: Exception received while running file "*scriptName.jacl*"; exception information: com.ibm.bsf.BSFException: error while evaluating Jacl expression: missing close-bracket
- Unexpected error CWSIV0806E in WebSphere log following deletion of an outbound service
- Separator exception
- Enabling authentication in the file transfer service
- The format of "\$AdminConfig list" output changed for V6.0
- You are not prompted for user ID and password after applying V6.0.2 service if you use an existing 6.0 profile
- When running the \$AdminApp searchJNDIReferences command with the Java Naming and Directory Interface (JNDI) name of a message destination, the message destination reference is not returned
- AWXJR0006E: The file, {0}, was not found.
- WASX7017E: Exception received while running file "<application stopping script>"; exception information: javax.management.MBeanException com.ibm.ws.exception.RuntimeWarning: Application <appname> not started

If you do not see your problem here:

- If you are not able to enter wsadmin command mode, try running **wsadmin -c "\$Help wsadmin"** for help in verifying that you are entering the command correctly.
- If you can get the wsadmin command prompt, enter **\$Help help** to verify that you are using specific commands correctly.
- wsadmin commands are a superset of Jacl (Java Command Language), which is in turn a Java-based implementation of the Tcl command language. For details on Jacl syntax beyond wsadmin commands, refer to the Tcl developers' site, <http://www.tcl.tk>. For specific details relating to the Java implementation of Tcl, refer to <http://www.tcl.tk/software/java>.
- Browse the *install_dir/profiles/profile_name/logs/wsadmin.traceout* file for clues.
 - Keep in mind that wsadmin.traceout is refreshed (existing log records are deleted) whenever a new wsadmin session is started.
 - If the error returned by wsadmin does not seem to apply to the command you entered, for example, you receive WASX7023E, stating that a connection could not be created to host "myhost," but you did not specify "-host myhost" on the command line, examine the properties files used by wsadmin to determine what properties are specified. If you do not know what properties files were loaded, look for the WASX7326I messages in the wsadmin.traceout file; there will be one of these messages for each properties file loaded.

WASX7016E, WASX7017E, or WASX7209I: Jython scripting language error

The following errors may occur when you run this Jython script:

Jython script

```
"profile_root/bin/wsadmin.sh -lang jython -profile profile_name -host host_name -f
script_file.py"
```

Error Messages

```
WASX7209I: Connected to process "server1" on node
node_name using SOAP connector; The type of process is:
UnManagedProcess
```

```
WASX7016E: Exception received while reading file
"script_file.py"; exception information:
sun.io.MalformedInputException
```

```
WASX7017E: Exception received while running file
"script_file.py"; exception information:
com.ibm.bsf.BSFException: exception from Jython: Traceback
(innermost last): File "<string>" line 89, in ? NameError: log
```

These errors can occur because there are UTF-8 characters in the file that are not valid. The default codepage for RHEL 3 is UTF-8 (en_US.UTF-8). When doing a text file read through Java™ code, the program assumes all characters are UTF-8 encoded. There might be one or more characters in the file that are not part of the UTF-8 specification, causing the load to fail. An easy way to determine if a character that is not valid is causing the error is to enter `export LANG=C` and run the script again. If you determine that the problem is a character that is not valid:

1. Open a new text reader on the file.
2. Read it one character at a time.
3. Print the character that is not valid.
4. When you press the back characters, you get the exception and will then know which of the characters is causing the error.
5. Remove any characters that are not valid, then run the script again

"WASX7023E: Error creating "SOAP" connection to host" or similar error trying to launch wsadmin command line utility

By default, the wsadmin utility attempts to connect to an application server at startup. This is because some commands act upon running application servers. This error indicates that no connection could be established.

To resolve this problem:

- If you are not sure whether an application server is running, start it by entering **startserver.sh server short name** from the command prompt. If the server is already running, you will see an error similar to "ADMU3027E: An instance of the server is already running".
- If you are running a z/OS configuration, you will first need to start the deployment manager by issuing the following command from a command prompt on the MVS console:

```
START dmgr_proc_name,JOBNAME=server_short_name,
ENV=cell_short_name.node_short_name.server_short_name
```

Note: This command must be entered on a single line. It is split here for display purposes.

Then you can launch wsadmin immediately to connect to the deployment manager, or start a node and application server to connect to.

- If an application server is running and you still get this error:
 - If you are running remotely (that is, on a different machine from the one running WebSphere Application Server), you must use the **-host hostname** option to the wsadmin command to direct wsadmin to the right physical server.
 - If you are using the **-host** option, try pinging the server machine from the command line from the machine on which you are trying to launch wsadmin to verify there are no issues of connectivity such as firewalls.

- verify that you are using the right port number to connect to the WebSphere Application Server process:
 - If you are not specifying a port number (using the `-port` option) when you start the `wsadmin` tool, the `wsadmin` tool uses the default port specified in `install_dir/profiles/profile_name/properties/wsadmin.properties` file, property name=`com.ibm.ws.scripting.port` (default value =8879).
 - The port that `wsadmin` should send on depends on the server process `wsadmin` is trying to connect to.

For a single-server installation, `wsadmin` attempts to connect to the application server process by default. To verify the port number:

- Look in the file `profile_root/config/cells/cell_name/nodes/node_name/serverindex.xml` for a tag containing the property `serverType="APPLICATION_SERVER"`.
- Look for an entry within that tag with the property `endpointName="SOAP_CONNECTOR_ADDRESS"`.
- Look for a `port` property within that tag. This is the port `wsadmin` should send on.

In a WebSphere Application Server, Network Deployment installation, `wsadmin` launched from the `bin` directory on the WebSphere Application Server, Network Deployment installation attempts to send requests to the deployment manager by default. To verify the port number:

- Get the hostname of the node on which the Deployment Manager is installed.
- Using that hostname, look in `profile_root/config/cells/cell_name/nodes/node_name/serverindex.xml` file for a tag containing the property `serverType="DEPLOYMENT_MANAGER"`.
- Within that tag, look for an entry with a property `endpointName="SOAP_CONNECTOR_ADDRESS"`.
- Within that tag, look for a "port" property. This is the port that the `wsadmin` tool should send on.
- If security is enabled, verify that the TSO or telnet userid that invokes the scripting client has a keyring by the name that is specified in the `ssl.client.props` file. The keyring must be correct to establish the SSL connection. The default name for the keyring is `WASKeyring`. This keyring contains the certificate authority certificate for the administrative server.

"com.ibm.bsf.BSFException: error while eval'ing Jacl expression: no such method *command name* in class com.ibm.ws.scripting.AdminConfigClient" returned from `wsadmin` command.

This error is usually caused by a misspelled command name. Use the `$AdminConfig help` command to get information about what commands are available. Note that command names are case-sensitive.

WASX7022E returned from running "`wsadmin -c ...`" command, indicating invalid command

If the command following `-c` appears to be valid, the problem may be caused by the shell attempting to do variable substitution. Variable substitution can occur on Unix System Services if `wsadmin -c` invokes a command that is enclosed in double quotes and includes dollar signs. To confirm that this is the problem, check the command to see if it contains an unescaped dollar sign, for example: `wsadmin -c "$AdminApp install"`.

To correct this problem, escape the dollar sign with a backslash. For example: `wsadmin -c "\$AdminApp install ..."`.

Note: When the command is enclosed in single quotes, the shell does not attempt to do variable substitution. Therefore, you do not need to escape the dollar sign. Example: `wsadmin.sh -c '$AdminApp install ...'`

com.ibm.ws.scripting.ScriptingException: WASX7025E: String "" is malformed; cannot create ObjectName

One possible cause of this error is that an empty string was specified for an object name. This can happen if you use one scripting statement to create an object name and the next statement to use that name, perhaps in an "invoke" or "getAttribute" command, but you don't check to see if the first statement really returned an object name. For example (the following samples use basic Jacl commands in addition to the wsadmin Jacl extensions to make a sample script):

```
#let's misspell "Server"
set serverName [$AdminControl queryNames type=Srever,*]
$AdminControl getAttributes $serverName
```

To correct this error, make sure that object name strings have values before using them. For example:

```
set serverName[$AdminControl queryNames node=mynode,type=Server,name=server1,*]
if {$serverName == ""} {puts "queryNames returned empty - check query argument"}
else {$AdminControl getAttributes $serverName}
```

For details on Jacl syntax beyond wsadmin commands, refer to the Tcl developers' site, <http://www.tcl.tk>.

WASX701E: Exception received while running file "*scriptName.jacl*"; exception information: com.ibm.bsf.BSFException: error while evaluating Jacl expression: missing close-bracket

This error is caused by a mix-up between the code page that the scripting client expects to see and the code page in which the Jacl script was written.

To fix this problem, set the `-Dscript.encoding=script codepage` option in the `wsadmin.sh` or `wsadmin.bat` file to the code page of the Jacl script. The following guideline will help you to determine the code page of the script:

- If the script was written in the OMVS interface using the OEDIT editor, the code page is IBM-037. In this case, set the option to the following: `-Dscript.encoding=Cp037`
- If the script was written in a telnet session to the OMVS interface using the VI editor, the code page is IBM-1047. In this case, set the option to the following: `-Dscript.encoding=Cp1047`
- IF the script was written on a personal computer, or any other ASCII machine, and was transferred to the host as a text file, the code page is IBM-1047. In this case, set the option to the following: `-Dscript.encoding=Cp1047`
- If the script was written on a personal computer, or any other ASCII machine, and transferred to the host in binary format, the code page is ISO-8859-1 (ASCII). In this case, you do not need to set the option because the default is ASCII. You should review other possible reasons for this error.

Unexpected error CWSIV0806E in WebSphere log following deletion of an outbound service

This error occurs when an exception is issued for destination `MPOutBoundServicePortDestination`, on messaging engine `trueliesNode01.server1-FVTSIBus01`, on bus `FVTSIBus01`, for endpoint activation:

```
com.ibm.websphere.sib.exception.SINotPossibleInCurrentConfigurationException: CWSIP0111E: The destination with name MPOutBoundServicePortDestination is being deleted on messaging engine {1}.
```

You can ignore this error; it is benign.

Separator exception

You must use forward slashes (/) as your path separator. Backward slashes (\) will not work.

Enabling authentication in the file transfer service

The file transfer service provides role-based authentication. Two versions of the file transfer web application are provided. By default, the version that does not authenticate its caller is installed. This default supports compatibility between the WebSphere Application Server, Network Deployment, Version 5.0 and Version 5.0.1 or later. Turning the file transfer authentication on is recommended to prevent unauthorized use of the file transfer application. However, if you have any Version 5.0 clients in your WebSphere Application Server, Network Deployment environment, they will not be able to communicate with the secured file transfer application if global security is turned on.

In WebSphere Application Server Version 6.x, mixed cells are supported and file transfer has become a system application. If all of the nodes in the cell are of Version 5.0.1 or later, you can activate authentication in the file transfer service by redeploying the file transfer application at the deployment manager. The compatible version is shipped in the `${app_server_root}/systemApps/filetransfer.ear` directory. The secured version is provided in the `${app_server_root}/systemApps/filetransferSecured.ear` directory.

Also in WebSphere Application Server Version 6.x, the file transfer service is also supported for WebSphere Application Server, Express as long as the node is not federated into a managed cell. If this node becomes federated, it will make use of the deployment manager file transfer.

A wsadmin Jacl script is provided to help you redeploy file transfer. The script is `redeployFileTransfer.jacl` and you can find it in the `${app_server_root}/bin` directory. After the deployment manager and all the nodes are at Version 5.0.1 or later, you can deploy the secured file transfer service by running the script. The syntax for running the script from the bin directory includes the following:

```
wsadmin -profile redeployFileTransfer.jacl -c "fileTransferAuthenticationXxx  
cell_name node_name server_name
```

where "Xxx" is "On" or "Off".

For example, when running the script to enable use of the `filetransferSecured.ear`, the syntax is similar to the following example:

```
wsadmin -profile redeployFileTransfer.jacl -c  
"fileTransferAuthenticationOn managedCell managedCellManager dmgr"
```

or

```
wsadmin -profile redeployFileTransfer.jacl -c  
"fileTransferAuthenticationOn baseCell base server1"
```

If you want to go back to run the file transfer service without authentication, you can run the script as shown in the following example:

```
wsadmin -profile redeployFileTransfer.jacl -c  
"fileTransferAuthenticationOff managedCell managedCellManager dmgr"
```

or

```
wsadmin -profile redeployFileTransfer.jacl -c  
"fileTransferAuthenticationOff baseNodeCell baseNode server1"
```

The format of "\$AdminConfig list" output changed in V6.0

If you have a script that parses the output of `$AdminConfig list`, such as `$AdminConfig list Node`, you might receive errors, such as "Node not found." Scripts should not parse the output of `$AdminConfig`; however, if you have a script that does this parsing, it must be updated for WebSphere Application Server V6.0 to reflect changes to the output format.

You are not prompted for user ID and password after applying V6.0.2 service if you use an existing 6.0 profile

If security is enabled, executing a .bat file requires a user ID and password. On V6.0.2, a new feature is introduced to prompt you for a user ID and password if they are not supplied in the command line. However, this feature is not available for profiles that were created at the 6.0 level.

Property files for profiles created at the V6.0 level are not updated after applying the V6.0.2 refresh pack.

There are two solutions to this problem:

1. Create a new profile after applying the V6.0.2 service. This new profile contains all the updated property files and you will then be prompted for a user ID and password.
2. If you want to keep the existing V6.0 profile and use the new prompt feature, you must manually update three files:
 - for `app_server_root/properties/soap.client.props`, add the following line:
`com.ibm.SOAP.loginSource=prompt`
 - for `app_server_root/properties/wsjaas_client.conf`, add the following lines:

```
WSAdminClientLogin {
    com.ibm.ws.security.common.auth.module.proxy.WSLoginModuleProxy required del
egate=com.ibm.ws.security.common.auth.module.WSAdminClientLoginModuleImpl;
};
```
 - for `app_server_root/bin/setupCmdLine.bat` add the following line:
`SET JAASSOAP=-Djava.security.auth.login.config=app_server_root/properties/wsjaas_client.conf`

When running the \$AdminApp searchJNDIReferences command with the Java Naming and Directory Interface (JNDI) name of a message destination, the message destination reference is not returned

This problem occurs when the command \$AdmnApp searchJNDIReferences is run with the JNDI name of a message destination. The command cannot collect the message destination reference that is defined in the application deployment descriptor. The message destination that you configured for the application server is defined with a message destination link on not one element, but two: both a message-driven bean (MDB) and a message destination reference.

Currently there is no workaround for this problem. The \$AdmnApp searchJNDIReferences command cannot return a reference for a message destination that is defined on two elements.

AWXJR0006E: The file, {0}, was not found.

This problem occurs when you attempt to configure a JACC provider for the Tivoli Access Manager using the wsadmin tool in a deployment manager environment with or without nodes added. You enter a deployment manager node name, for example, t54Manager, instead of an asterisk (*) for all nodes. The wsadmin command finishes successfully but when you try to add a new node and start the node, or start an existing node, you receive an error in the nodeagent SystemOut.log file similar to the following:

```
[12/7/05 17:09:51:266 CST] 0000000a SystemOut      0 AWXJR0006E  The file,
C:\cc_was602\WebSphere\AppServer\profiles\AppSrv01\etc\tam\amwas.t54Node01_.amjacc.pr
operties, was not found.
[12/7/05 17:09:51:266 CST] 0000000a distSecurityC E   SECJ0391E: Error when setting
the Policy object to the providers policy implementation {0}. The exception is
{1}.
[12/7/05 17:09:51:281 CST] 0000000a distSecurityC E   SECJ0324E: Error during Java
2 Security and Dynamic Policy initialization.
```

To workaround this problem, unconfigure the existing Tivoli Access Manager configuration and configure it again using an asterisk (*) for the node name, for example:


```
wsadmin.bat -user wsadmin -password pw1 -f enableTAM.jacl "*" TAMHostName:7135"  
TAMHostName:7136:1" "cn=wsadmin,o=ibm,c=us" "o=ibm,c=us" "sec_master"  
sec_master pw1 "9990:9999"
```

WASX7017E: Exception received while running file "<application stopping script>"; exception information: javax.management.MBeanException com.ibm.ws.exception.RuntimeWarning: Application <apname> not started

While using scripting, this error is issued when you stop an application that is already stopped or start an application that is already running.

Tracing and logging facilities - troubleshooting tips

Use this information if you are having problems using tracing, logging or other troubleshooting tools.

What kind of problem are you having?

- Netscape browser fails when trying to enable a component trace.

Netscape browser fails when trying to enable a component trace

On systems using AIX, the Netscape browser fails when you try to enable trace on a component.

To work around this problem, do one of the following:

- Disable JavaScript on the browser and continue setting trace.
- Administer the AIX server from a remote machine running another browser and operating system.
- Change the trace manually in the *server.xml* file.

Cannot restart the Deployment Manager monitoring policy

The Deployment Manager monitoring policy is permanently set to STOPPED and cannot be changed. Therefore, if the Deployment Manager fails it will never be restarted by WebSphere Application Server monitoring.

To correct the problem, use the z/OS Automatic Restart Facility (ARM) to monitor and restart the Deployment Manager. See z/OS Automatic Restart Facility for information on ARM.

Timeout condition resolutions

This topic gives an overview of how to resolve timeout conditions.

In such a complex environment as WebSphere Application Server for z/OS, timeouts might occur for many different reasons. Although you can alter timeout values, you should not do so until you understand why the timeout occurs. Depending on the timeout condition, you might be able to permanently fix the timeout condition by doing some system or application tuning. For example, if the diagnostic data indicates throughput problems, you can alter the number of server regions, the number of threads within each server region, or the use of replicated servers.

Generally speaking, increasing the timeout values should be your last resort, or only a temporary action taken to prevent multiple timeout-abend dumps from causing system performance problems. If you increase timeout values without properly diagnosing the timeout condition, the only results you might see are less frequent abends and dumps for the same timeout condition, or slower system or application performance.

Timer overview

Timers define a limit to the amount of time that is required for a specific operation to complete. The type of operation a timer controls determines when the time period, that is set for that timer, starts to elapse.

Timer properties as they relate to configuring your message-driven beans to work with listener ports or activation specifications.

For WebSphere Application Server Version 7 and later, listener ports are deprecated. Therefore you should plan to migrate your WebSphere MQ message-driven bean deployment configurations from using listener ports to using activation specifications. However, you should not begin this migration until you are sure the application does not have to work on application servers earlier than WebSphere Application Server Version 7. In some cases you will continue to use the WebSphere MQ message-driven bean deployment and listener ports and in other case you will use the WebSphere MQ message-driven bean deployment and activation specifications.

The following properties DO NOT apply to activation specification driven message-driven bean deployment. That is, the properties require you to configure them to use the WebSphere MQ message-driven bean deployment and listener ports:

- control_region_mdb_request_timeout
- control_region_mdb_queue_timeout_percent
- server_region_mdb_stalled_thread_dump_action

The follow properties DO apply to activation specification driven message-bean deployment. That is, these properties require you to configure them to use the WebSphere MQ message-driven bean deployment and activation specifications.

- control_region_wlm_dispatch_timeout
- control_region_iiop_queue_timeout_percent
- server_region_iiop_stalled_thread_dump_action

As you follow the instructions to configure these properties, remember what properties apply to listener ports versus activation specifications.

Most timers have a default value that defines a reasonable time period, during which a particular operation should complete. When the time limit specified for the timer elapses, the product takes one of the following actions:

- Sends a minor code to the client if the time limit elapses before the client request is dispatched to a servant.
- Abnormally ends the servant with an EC3 ABEND if the time limit elapses while the client request is being processed by an application component running in the servant.

In this situation all of the threads in this servant are terminated. The servant is also terminated to prevent the application from holding resources, and thereby preventing other requests from using these resources. After the servant is terminated, workload management (WLM) starts a new servant to take the place of the servant that was terminated.

gotcha: The total transaction lifetime timeout and the maximum transaction timeout have grace periods beyond the timeout value specified of approximately four minutes. This extended time period must elapse before an ABEND occurs.

Different types of timers might reach their time limits simultaneously, because the operations they control might overlap to a certain degree. For example, suppose the application server receives an IIOp client request that is processed by an application component that uses transaction support. In this case, both of the following timers count down simultaneously:

- control_region_wlm_dispatch_timeout, which limits both the amount of time a client request waits on the WLM queue, as well as the time required for the application component to process the request
- transaction_defaultTimeout, which limits the amount of time the controller waits for a transaction to be either committed or rolled back

These timers overlap only for the time during which the transaction is processed. To determine which timer cause the error, you can use the symptoms- specific minor codes or EC3 abend reason codes.

To determine the occurrence of a timeout as quickly as possible and to prevent further resource locking, WebSphere Application Server prevents further transactional work on the transactional path where the timeout condition has taken place. This applies equally to attempting to perform work under the current transaction context and to attempting to perform work under a different transactional context.

The timers that are used to control processing behavior can be classified into five general types. These general types, and the operations that they control, are summarized in the following table.

Table 72. Summary of general types. The timers that are used to control processing behavior can be classified into five general types. These general types, and the operations that they control, are summarized in the following table.

General type	Timer processing	Timeout symptoms
Input	Input timers define limits for receiving a complete request; the countdown starts when a connection to the J2EE server is established. The communication protocol (HTTP, HTTPS) determines the timer used for the request.	The session is closed.
Session	Session timers define limits for the use of session connections. These timers start the countdown as soon as a session becomes idle.	The session is closed.
WLM dispatch	Dispatch timers control how long a complete client request waits to be dispatched in a servant region for processing. For some dispatch timers, an additional value can be specified that designates the percent of the dispatch time as the timeout value for the WLM queue. If this time is exceeded, then the work is removed from the WLM queue, but an abend is not issued for the servant. The countdown starts when the controller places the request on the WLM queue. Depending on the specific timer, the time limit can include both the wait time on the WLM queue, and the time required for processing a response to the client request.	Message BBOO0327I for all timeouts. If the servant is terminated, then Message BBOO0232W and an EC3 ABEND are displayed in the servant, with one of the following reason codes: <ul style="list-style-type: none"> • 04130003 • 04130004 • 04130006
Transaction	Transaction timers define how long: <ul style="list-style-type: none"> • An application or controller waits for one transaction to complete. The countdown starts when the container starts a transaction on behalf of the application component. • A controller attempts to recover in-doubt transactions during peer restart and recovery mode. 	Message BBOT0003W or BBOO0232W, and an EC3 ABEND are displayed in the servant, with one of the following reason codes: <ul style="list-style-type: none"> • 04130002 • 04130005
Output	Output timers define how long a controller waits to receive output for a client request. The countdown starts when the client request is dispatched to the servant for processing. The communication protocol (HTTP or HTTPS) determines the timer used for the request.	Message BBOO0232W and an EC3 ABEND are displayed in the servant, with reason code 04130007.

Timeout conditions: analyzing diagnostic data

The following guidelines provide instructions for finding diagnostic data in an SVC dump that can help you determine what timeout condition occurred.

You should start by finding the task with the EC3 abend:

1. Format the TCB summary for the servant that was timed out by entering the following command:

```
ip summ format asid(x' address ')
```

where *address* is the address space ID of the servant.

Find the TCB that had the EC3 completion code. Ignore any EC3 completion code on the "main" thread which is the 4th TCB listed in the summary format (the 1st one after the 3 MVS TCBs). The WebSphere main thread is the one that is waiting in BBO_BOA::impl_is_ready. No application requests are ever dispatched on this thread, therefore there is nothing to timeout. During timeout processing the main thread for the server region is also abended with EC3 as a mechanism of bringing the address space down. Thus the reason why the EC3 completion code may appear on the main thread. This is never the cause of a timeout though, only a result of timeout processing.

2. If there is no EC3 completion code in the TCB summary, look in systrace. Format the systrace in GMT time since the other timestamps you'll be comparing it to are in GMT time. To format in GMT time, enter the following command::

```
ip systrace all time(gmt)
```

You may not see the EC3 abend in systrace either as systrace can cover a small amount of time.

3. You can also try looking in ip verbx mtrace or in syslog to see when the EC3 abend occurred. You'll need this time to determine the 'end' time of the request which is the GMT time the timeout value was reached.

You can determine what timeout values are in effect by checking the reason code associated with the EC3 abend.

Table 73. Reason code and explanations. You can determine what timeout values are in effect by checking the reason code associated with the EC3 abend.

Reason code	Explanation
04130002	The controller issued an ABTERM for this servant region because a transaction timeout occurred. Code under dispatch could have been in a tight loop.
04130003	The controller issued an ABTERM for this servant region because it was hung trying to move a controller request into the servant region. The target request was timed out, but the servant was currently copying the request. The controller checked the servant for progress at regular intervals, before taking action by issuing an ABTERM.
04130004	The controller issued a ABTERM for this servant region because the WLM queue timeout occurred. Code under dispatch could have been in a tight loop.
04130005	The controller issued an ABTERM for this servant region because a transaction timeout occurred. The transaction has timed out, but no current request associated with the transaction was found. The servant associated with the transaction will be terminated.
04130006	A controller thread encountered a problem while processing a request. The request has been queued to WLM and associated with a servant region. The termination of the associated servant region is needed to complete cleanup for the request.
04130007	The controller issued a ABTERM for this servant region because the HTTP OUTPUT timeout occurred. Code under dispatch could have been in a tight loop.

You can try to find the method name to determine if it was

```
httpRequest
```

,
httpsRequest

or

DispatchbyURI

or some other method. If the request is not specifically a request that came through the HTTP or HTTPS transport handlers, the

protocol_http_output_timeout

(HTTP) and

protocol_https_timeout_output

(HTTPS) timeout values will not be a factor. In other words, when the request is a

DispatchbyURI

method, the request is received through the RMI/IIOP protocol, so the

protocol_http

* variables have no affect.

You can then use the IPCS verbexit LEDATA, with the CEEDUMP, or NTHREADS option to obtain the stack trace for that request.

Timeout conditions - possible causes and fixes

This file lists common timer variables and tools for monitoring these timeout conditions

The timer that expires first might not indicate the actual problem that needs to be fixed. To properly diagnose timeout conditions, you should know all of the timer values that might be in effect for a particular servant region.

Table 74. Timeout conditions - possible causes and fixes. You can use these common timer variables to monitor timeout conditions.

General type of timer	Possible causes	Possible solutions
Input	The client sent only part of the data and was delayed in sending the rest.	The application on the client side may want to consider having retry logic in place if it does receive a timeout minor code in return.
Session	The session is idle through lack of use.	If you consider losing idle sessions to be a problem, increase the values of the persistent-session timeouts, or use the session more frequently.

Table 74. Timeout conditions - possible causes and fixes (continued). You can use these common timer variables to monitor timeout conditions.

General type of timer	Possible causes	Possible solutions
WLM dispatch	<p>No threads are free to pick up the request because of one of the following conditions:</p> <ul style="list-style-type: none"> • The threads are all busy processing requests. • The currently executing threads are waiting for a response from DB2, WebSphere MQ, another server, and so on. In this case, look for messages indicating contention for resources; for example, on the z/OS console, you might see messages about DB2 deadlocks. <p>In either case, the request times out waiting in the WLM queue to be dispatched in a servant (region).</p>	<p>The case where the threads are all busy processing requests could indicate one of the following conditions:</p> <ul style="list-style-type: none"> • The number of servant regions that WLM may start is set too low. To set this value, on the administrative console select Servers >Application Servers >server_name >Server Instance. Click on Multiple Instances Enabled and specify a value for Maximum Number of Instances. • The number of threads allowed within a servant region is set too low. The number is controlled by the Isolation Policy setting in Administrative console or WebSphere variable: <i>server_region_workload_profile</i> • You need to replicate servers to handle the amount of incoming work. <p>All of these conditions indicate that performance tuning might be necessary.</p>
Transaction	<p>Possible causes of transaction timeouts include:</p> <ul style="list-style-type: none"> • The same as those for WLM dispatch timeouts, or • Delays that prevent the transaction coordinator from committing or rolling back a transaction within the allotted time. 	<p>See the possible solutions for WLM dispatch timeouts. In addition, you can look for messages indicating contention for resources that are involved in the transaction that timed out.</p>
Output	<p>Possible causes of output timeouts are the same as those for WLM dispatch (dispatch is for IIOp, output is for HTTP).</p>	<p>See the possible solutions for WLM dispatch timeouts. In addition, you can use the WebSphere variable <i>protocol_accept_http_work_after_min_srs=1</i> to prevent the HTTP transport handler from dispatching requests until WLM starts a minimum number of servant regions.</p>

Timeout values: guidelines for altering timeout values

This file lists common timer variables and tools for monitoring these timeout conditions

Generally speaking, increasing the timeout values should be your last resort, or only a temporary action taken to prevent multiple timeout-abend dumps from causing system performance problems. If you increase timeout values without properly diagnosing the timeout condition, the only results you might see are less frequent abends and dumps for the same timeout condition, or slower system or application performance.

For information on how to set values for these timer variables, and how these variables map to internal variables, see Controlling behavior through timeout values

(Some WebSphere variables are split on multiple lines for printing purposes.)

Table 75. Common timer variables and tools for monitoring these timeout conditions. Ensure that you diagnose the timeout condition before increasing timeout values.

WebSphere variable and its relationship, if any, to other timers	How to monitor processing for this type of timeout condition:	If you want to adjust the value, consider the following:
<p>WLM timeout</p> <p>For HTTP work and Scalable Messaging Support, the WLM timer is not set and only the ConnectionResponseTimeout is in effect (covering the entire dispatch window)</p>	SMF provides data on WLM queue time	How long work takes to get to a servant depends on the number of servants that WLM starts, how many you let it start, how many service classes the work is spread across, how much work you're getting, and so on.
<p>ConnectionIOTimeout</p> <p>None.</p>	This behavior is not easily monitored. Turning on a trace point would indicate whether a client failed because of this input timeout setting, but tracing has performance consequences.	<ul style="list-style-type: none"> How long are you willing to allow a control region worker thread to be blocked while it is waiting for a message? How big are incoming HTTP requests? The larger they are, the longer it might take to get the whole request through the network.
<p>ConnectionResponseTimeout</p> <p>If the application component starts transactions, then the transaction timers also might be involved.</p>	This behavior is not easily monitored, but the controller will terminate the servant (region) with abend EC3 for this timeout condition.	<ul style="list-style-type: none"> How long are you willing to let a client hang waiting for a response? How long are you willing to let a thread in a servant (region) be tied up working on a response before concluding that the request has taken too long? If you have multiple application threads in the servant (region), all of them will be terminated when only one of them times out. This loss of work might make you want to allow these time outs to occur less frequently.
<p>ConnectionKeepAliveTimeout</p> <p>None. All the other timers relate to work processing, whereas this one relates to what happens when there is no work.</p>	None.	How much time passes between requests vs. how much does it cost to establish a new session. You would want to keep idle sessions around for a while to avoid the startup cost of a new session, but don't want to keep them forever as resource usage accumulation will eventually be a problem.
<p>Request Timeout (ORB Service)</p> <p>None. This variable is a client-side timeout, and IOP only.</p>	None, other than to observe the timeouts occurring on the client side.	How long are you willing to let the client wait?
<p>ORB listener keep alive ORB SSL listener keep alive</p> <p>None. These variables relate to session activity during idle periods and only for IOP, so these timers do not interact with the ConnectionKeepAliveTimeout timer.</p>	You should read TCP/IP APAR PQ18618 for information about the SOCK_TCP_KEEPLIVE values and their consequences.	Is it useful to have idle sessions timeout? They normally don't which can consume resources. However, detecting a timeout requires network traffic between TCP/IP stacks. Creating traffic on otherwise idle sessions may have network consequences you don't want.
<p>Total Transaction Lifetime Timeout</p> <p>This variable can be overridden by applications up to the maximum indicated by the Maximum Transaction Timeout variable, which limits the amount of time an application can set for its transactions to complete. Output timers also might cause work to time out, but the transaction timers and output timers are not aware of each other.</p>	The controller issues message BBOT0003W to indicate a timeout condition, and terminates the servant (region) with abend EC3 reason codes 04130002 or 04130005.	<ul style="list-style-type: none"> How long are you willing to let a client hang waiting for a response? How long are you willing to let a thread in a servant (region) be tied up working on a response before concluding that the request has taken too long? If you have multiple application threads in the servant (region), all of them will be terminated when only one of them times out. This loss of work might make you want to allow these time-outs to occur less frequently.
<p>Maximum Transaction Timeout</p> <p>If set, this variable limits the amount of time an application can set for its transactions to complete. If the Maximum Transaction Timeout variable is not set, application transactions are controlled by the time limit set on the Total Transaction Lifetime Timeout variable.</p>	None.	Same considerations as for transaction_defaultTimeout
<p>transaction_recoveryTimeout</p> <p>None</p>	None.	Locks are held while one controller waits for other controllers that are required to resolve in-doubt transactions. How long can you afford to have these resources held?

Table 75. Common timer variables and tools for monitoring these timeout conditions (continued). Ensure that you diagnose the timeout condition before increasing timeout values.

WebSphere variable and its relationship, if any, to other timers	How to monitor processing for this type of timeout condition:	If you want to adjust the value, consider the following:
<code>server_region_request_cputimeused_limit</code>	This behavior is not easily monitored, but the controller terminates a request when the specified cpu use time limit is reached.	<ul style="list-style-type: none"> • How much of your CPU time are you willing to let a single application request consume before taking some action against that request? • If you have multiple application threads in the servant, all of them are terminated if terminating this request means that the servant now has enough unresponsive threads for the controller to terminate the servant. The setting for the <code>server_region_stalled_thread_threshold_percent</code> property determines how many threads need to be unresponsive before the servant is terminated.
<code>server_region_stalled_thread_threshold_percent</code>	This behavior is not easily monitored, but the controller will terminate the servant with abend EC3 when the percentage of unresponsive threads meets this condition.	<ul style="list-style-type: none"> • How long are you willing to let a client wait for a response? • How long are you willing to let a thread in a servant be tied up working on a response before concluding that the request has taken too long? • If you have multiple application threads in the servant, all of them are terminated when the controller terminates the servant because the percentage of unresponsive threads is reached. This loss of work might make you want to allow a larger percentage of request threads to become unresponsive before having the controller terminate the servant.

Appendix. Directory conventions

References in product information to *app_server_root*, *profile_root*, and other directories imply specific default directory locations. This topic describes the conventions in use for WebSphere Application Server.

Default product locations - z/OS

app_server_root

Refers to the top directory for a WebSphere Application Server node.

The node may be of any type—application server, deployment manager, or unmanaged for example. Each node has its own *app_server_root*. Corresponding product variables are *was.install.root* and *WAS_HOME*.

The default varies based on node type. Common defaults are *configuration_root*/AppServer and *configuration_root*/DeploymentManager.

configuration_root

Refers to the mount point for the configuration file system (formerly, the configuration HFS) in WebSphere Application Server for z/OS.

The *configuration_root* contains the various *app_server_root* directories and certain symbolic links associated with them. Each different node type under the *configuration_root* requires its own cataloged procedures under z/OS.

The default is */wasv8config/cell_name/node_name*.

plug-ins_root

Refers to the installation root directory for Web Server Plug-ins.

profile_root

Refers to the home directory for a particular instantiated WebSphere Application Server profile.

Corresponding product variables are *server.root* and *user.install.root*.

In general, this is the same as *app_server_root/profiles/profile_name*. On z/OS, this will always be *app_server_root/profiles/default* because only the profile name "default" is used in WebSphere Application Server for z/OS.

smpe_root

Refers to the root directory for product code installed with SMP/E or IBM Installation Manager.

The corresponding product variable is *smpe.install.root*.

The default is */usr/lpp/zWebSphere/V8R0*.

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.

APACHE INFORMATION. This information may include all or portions of information which IBM obtained under the terms and conditions of the Apache License Version 2.0, January 2004. The information may also consist of voluntary contributions made by many individuals to the Apache Software Foundation. For more information on the Apache Software Foundation, please see <http://www.apache.org>. You may obtain a copy of the Apache License at <http://www.apache.org/licenses/LICENSE-2.0>.

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 Intellectual Property & Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
USA

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation
Mail Station P300
2455 South Road
Poughkeepsie, NY 12601-5400
USA
Attention: Information Requests

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

Trademarks and service marks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. For a current list of IBM trademarks, visit the IBM Copyright and trademark information Web site (www.ibm.com/legal/copytrade.shtml).

Microsoft and Windows 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.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Other company, product, or service names may be trademarks or service marks of others.

Index

A

- administrative console 18
 - accessing help 29
 - buttons 5
 - command assistance 30
 - actions 31
 - features 9
 - logging off 20
 - product installation 19
 - product uninstallation 19
 - saving changes 22
 - settings 23
 - preferences 26
 - scopes 27
 - specifying preferences 23
 - starting 20
 - usage 5
- Ant 187
- application files
 - adding 333
 - changing 333
- applications
 - administering
 - targets 205
 - AppManagement MBean 383
 - bindings 284
 - common deployment framework 206
 - configuration
 - binary location usage 295
 - enterprise applications 283
 - deployment
 - targets 205
 - deployment descriptors 312
 - deployment targets 322
 - disabling automatic start 318
 - dynamic reloading 330
 - editing applications 362
 - enterprise applications 290
 - export copy 344
 - export data 345
 - exporting from servers 342
 - hot deployment 330
 - installation bindings 223
 - management 356
 - managing 351, 352
 - partial applications
 - adding to 366
 - deleting 366
 - updating 366
 - removing application files 346
 - resolving configuration conflicts 340
 - session sharing 359
 - settings 292
 - binary 297
 - installation 222
 - startup behavior 294
 - updates 326

- applications (*continued*)
 - starting 316
 - startup configuration 293
 - status 319
 - stopping 316
 - system applications 206
 - uninstallation 347
 - administrative console 345
 - updating
 - administrative console 324
 - updating files 320
- assets 398
 - deleting 403
 - exporting 404
 - importing 392
 - managing 398
 - settings 394
 - adding 411
 - updating 401
 - uploading 394
 - updating 399
- authentication
 - JASPI enablement
 - applications 251

B

- backend ID
 - settings 271
- business-level applications 388, 408
 - administering 469
 - assets 390
 - composition units 391
 - creating 404, 420
 - administrative console 406
 - deleting 466
 - programming 471
 - deleting 548
 - editing 538
 - listing 509
 - starting 498
 - stopping 502
 - SCA 423
 - administrative console 425, 455
 - settings 409, 415
 - starting 457
 - stopping 458
 - updating 459

C

- class loaders 199, 301
 - Java EE 193
 - Java EE applications 200
 - settings 200
 - web modules 202
 - WebSphere server 198

- class loading
 - settings
 - update detection 303
- client modules
 - property settings 241
 - settings 241
- command-line tools 110
- commands
 - addNode 121
 - best practices 127
 - cleanupNode 131
 - EARExpander 155
 - managesdk 149
 - registerNode 134
 - renameNode 132
 - START 164
 - startManager 115
 - startNode 117
 - STOP 165
 - stopNode 118
 - syncNode 131
- composition units
 - settings 417, 433
 - adding 411

D

- deployment
 - applications
 - targets 205
 - business-level applications 387
 - EAR files
 - default bindings 220
 - enterprise modules with JSR-88 280
 - Java EE application files
 - WebSphere targets 207
 - Java EE files
 - administrative console 213
 - Java EE modules
 - deployment targets 208, 210
- directory
 - installation
 - conventions 605

E

- EJB Deploy
 - settings 260
- EJB JAR files
 - adding 337
 - changing 337
- Enterprise JavaBeans (EJB)
 - JNDI names for beans 243
 - references 246
 - settings
 - binding EJB business settings 244
- environment entries
 - settings
 - applications 269
 - client module 267
 - EJB modules 268

H

- HTTP plug-ins
 - configuration changes 339

I

- installation options
 - settings 229

J

- Java EE modules
 - mapping
 - to WebSphere servers 307
- JMS
 - binding settings
 - SCA composites 437
- JNDI
 - settings
 - JCA objects 272
- job managers
 - commands
 - restoreJobManager 163

M

- messages
 - settings
 - destination references 270
- metadata
 - settings
 - modules 275, 315
- module build ID
 - settings 278
- module customization
 - DConfigBeans 282
- modules
 - settings 238, 305

O

- options
 - relationship options
 - settings 414
 - settings 412

P

- policy sets
 - settings
 - attaching 429
- policy sets and bindings
 - references 446
 - service provider 443
- programming
 - adding assets 534
 - adding composition units 528
 - adding files 373
 - adding modules 368
 - application installation 353

- programming (*continued*)
 - application uninstallation 376
 - asset deletion 491
 - attribute manipulation 360
 - business-level applications 469, 471
 - status checking 505
 - composition unit deletion 551
 - deleting business-level applications 548
 - editing assets 487
 - editing business-level applications 538
 - editing composition units 542
 - exporting assets 495
 - file deletion 380
 - file updates 375
 - importing assets 475
 - listing assets 480
 - listing business-level applications 509
 - listing composition units 513
 - listing control operations 516
 - module deletion 378
 - module updates 370
 - starting applications 358
 - starting business-level applications 498
 - stopping business-level applications 502
 - updating applications 364
 - viewing assets 483
 - viewing business-level applications 520
 - viewing composition units 524

R

- roles
 - RunAs
 - SCA composites 433

S

- SCA
 - business-level applications
 - administrative console 455
 - composite artifacts
 - updating 460
 - composite definitions 461
 - deployment 421
 - domain information 462
 - Java EE composition unit relationships 428
 - JMS bindings
 - references and services 463
 - settings
 - component references 442
 - component services 443

- SCA (*continued*)
 - settings (*continued*)
 - composite components 441
 - HTTP endpoint URLs 441
 - service clients 452
 - service provider 448
 - virtual hosts 427
 - security
 - command-line tools 164
 - SCA
 - mapping roles to users 431
 - shared library reference and mapping
 - settings 262
 - shared library relationship and mapping
 - settings 263, 410

T

- target mapping
 - settings 413
- task automation
 - using Ant 187

V

- virtual hosts
 - mapping
 - for web modules 309
 - settings 249, 311

W

- WAR files
 - adding 335
 - changing 335
- web applications
 - settings
 - initial parameters for servlets 266
- web services
 - samples installation 221
 - settings
 - options to perform web services deployment 277
- WSDL
 - exporting documents 464

X

- XSD
 - exporting documents 464